



US012079626B2

(12) **United States Patent**
Brebner

(10) **Patent No.:** **US 12,079,626 B2**
(45) **Date of Patent:** ***Sep. 3, 2024**

(54) **METHODS AND SYSTEMS FOR CREATING APPLICATIONS USING SCENE TREES**

(71) Applicant: **Umajin Inc.**, Woburn, MA (US)
(72) Inventor: **David Brebner**, Palmerston North (NZ)
(73) Assignee: **Umajin Inc.**, Woburn, MA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/825,798**

(22) Filed: **Mar. 20, 2020**

(65) **Prior Publication Data**

US 2020/0218519 A1 Jul. 9, 2020

Related U.S. Application Data

(63) Continuation of application No. 16/047,613, filed on Jul. 27, 2018, now Pat. No. 10,599,405, which is a (Continued)

(51) **Int. Cl.**
G06F 8/30 (2018.01)
G06F 8/34 (2018.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 8/76** (2013.01); **G06F 8/311** (2013.01); **G06F 8/315** (2013.01); **G06F 8/34** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC G06F 8/34; G06F 8/36; G06F 8/38; G06F 8/41; G06F 8/60; G06F 8/76; G06F 8/311;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,673,331 A 9/1997 Lewis et al.
5,706,364 A 1/1998 Kopec et al.
(Continued)

FOREIGN PATENT DOCUMENTS

KR 20100055993 A 5/2010
WO 2016063022 A1 4/2016

OTHER PUBLICATIONS

Zhou Wenjun et al., Study on SaaS Architecture of Web3D Scene Content Management Platform and Method of Editing and Browsing, IEEE, 2012, retrieved online on Jun. 9, 2021, pp. 514-517. Retrieved from the Internet: <URL: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6272825>>. (Year: 2012).*

(Continued)

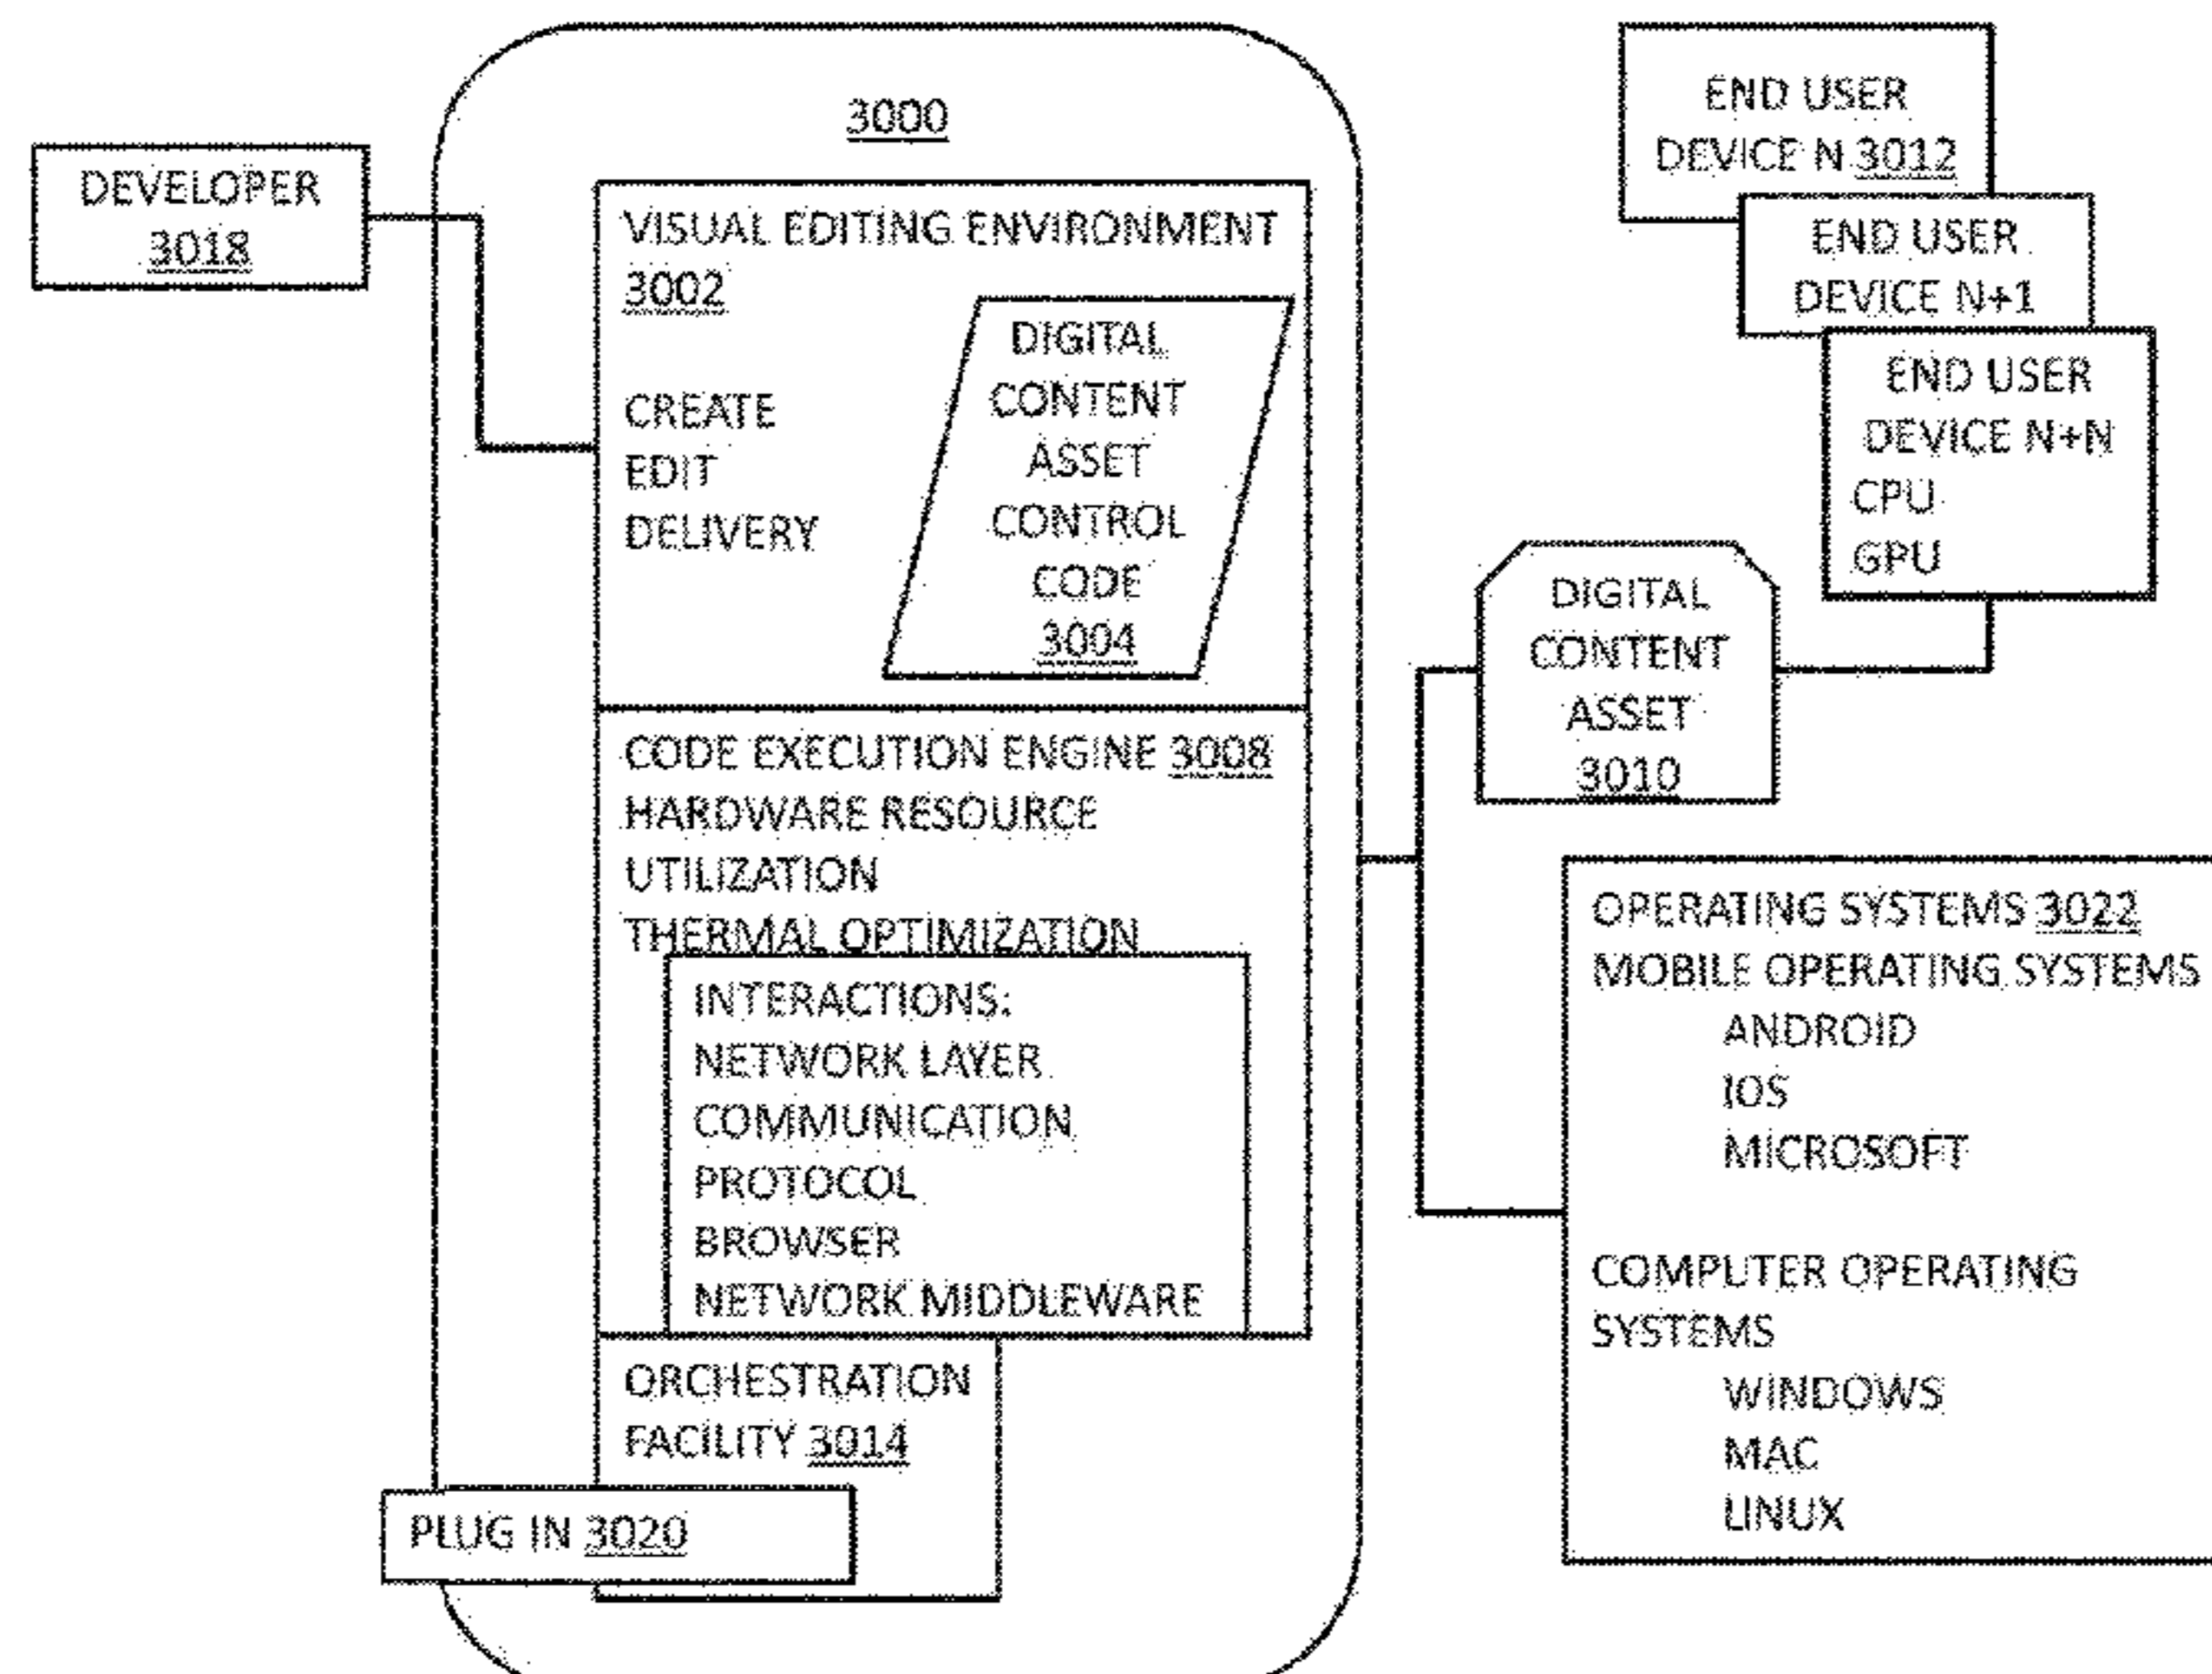
Primary Examiner — Hanh Thi-Minh Bui

(74) *Attorney, Agent, or Firm* — RMCK Law Group, PLC

(57) **ABSTRACT**

Systems and method for creating applications are disclosed. In embodiments, the system includes a visual editing environment configured to receive code that defines an application using a declarative language from one or more developers, wherein the code includes a scene tree description; a code execution engine configured to operate on the code created in the visual editing environment to create a scene tree that defines properties and behaviours of the application based on the scene tree description; and a publisher configured to publish the application, whereby instances of the application may be downloaded to respective client devices, wherein each instance of the application includes a respective instantiation of the scene tree, wherein the scene tree description is editable via the visual editing environment, such that edits to the scene tree description are updated to the respective instantiations of the scene tree at runtime.

22 Claims, 49 Drawing Sheets



Related U.S. Application Data					
	continuation of application No. PCT/US2018/035953, filed on Jun. 5, 2018.				
(60)	Provisional application No. 62/619,348, filed on Jan. 19, 2018, provisional application No. 62/559,940, filed on Sep. 18, 2017, provisional application No. 62/515,277, filed on Jun. 5, 2017.				
(51)	Int. Cl.				
	G06F 8/36 (2018.01)				
	G06F 8/38 (2018.01)				
	G06F 8/41 (2018.01)				
	G06F 8/60 (2018.01)				
	G06F 8/76 (2018.01)				
	G06F 9/451 (2018.01)				
	G06F 11/36 (2006.01)				
	H04L 67/02 (2022.01)				
	H04L 67/104 (2022.01)				
	H04W 4/80 (2018.01)				
(52)	U.S. Cl.				
	CPC G06F 8/36 (2013.01); G06F 8/38 (2013.01); G06F 8/41 (2013.01); G06F 8/60 (2013.01); G06F 9/451 (2018.02); G06F 11/362 (2013.01); H04L 67/02 (2013.01); H04L 67/104 (2013.01); H04W 4/80 (2018.02)				
(58)	Field of Classification Search				
	CPC G06F 8/315; G06F 9/44526; G06F 9/451; G06F 11/362; H04L 67/02; H04L 67/104; H04W 4/80				
	See application file for complete search history.				
(56)	References Cited				
	U.S. PATENT DOCUMENTS				
	5,760,788 A 6/1998 Chainini et al.	7,627,860 B2 12/2009 Kodosky et al.			
	5,764,241 A 6/1998 Elliott et al.	7,685,577 B2 3/2010 Pace et al.			
	5,911,070 A 6/1999 Solton et al.	7,698,398 B1 4/2010 Lai			
	6,047,129 A 4/2000 Frye	7,707,547 B2 4/2010 Colton et al.			
	6,145,120 A 11/2000 Highland	7,760,134 B2 7/2010 Morinaga et al.			
	6,167,563 A 12/2000 Fontana et al.	7,765,529 B1 7/2010 Singh et al.			
	6,282,699 B1 8/2001 Zhang et al.	7,865,888 B1 1/2011 Qureshi et al.			
	6,286,138 B1 9/2001 Purcell	7,917,888 B2 3/2011 Chong et al.			
	6,425,120 B1 7/2002 Morganelli et al.	7,971,194 B1 6/2011 Gilboa			
	6,438,594 B1 8/2002 Bowman-Amuah	7,974,825 B2 7/2011 Linebarger et al.			
	6,456,308 B1 9/2002 Agranat et al.	8,010,783 B1 8/2011 Cahill			
	6,496,850 B1 12/2002 Bowman-Amuah	8,019,725 B1 9/2011 Mulligan et al.			
	6,519,767 B1 2/2003 Carter et al.	8,060,400 B2 11/2011 Wellman			
	6,637,020 B1 10/2003 Hammond	8,060,864 B1 11/2011 Michelsen			
	6,654,950 B1 11/2003 Barnishan	8,064,928 B2 11/2011 Venkatachalam			
	6,823,495 B1 11/2004 Vedula et al.	8,239,824 B2 8/2012 Cifra et al.			
	6,836,884 B1 12/2004 Evans et al.	8,301,994 B1 10/2012 Shah et al.			
	7,020,662 B2 3/2006 Boreham et al.	8,347,276 B2 1/2013 Schadow			
	7,043,696 B2 5/2006 Santori et al.	8,352,903 B1 1/2013 Friedman			
	7,054,901 B2 5/2006 Shafer	8,417,728 B1 4/2013 Anders et al.			
	7,096,248 B2 8/2006 Masters et al.	8,468,535 B1 6/2013 Keagy et al.			
	7,149,770 B1 12/2006 Kalva et al.	8,677,318 B2 3/2014 Mohindra et al.			
	7,152,229 B2 12/2006 Chong et al.	8,688,676 B2 4/2014 Rush et al.			
	7,246,315 B1 7/2007 Andrieu et al.	8,713,530 B2 4/2014 Waite et al.			
	7,316,003 B1 1/2008 Dulepet et al.	8,726,264 B1 5/2014 Allen et al.			
	7,353,502 B2 4/2008 Stewart et al.	8,788,569 B2 7/2014 Griffiths et al.			
	7,382,248 B2 6/2008 Black et al.	8,825,750 B2 9/2014 Shimamoto			
	7,451,432 B2 11/2008 Shukla et al.	8,930,881 B1 1/2015 Eddins et al.			
	7,464,373 B1 12/2008 Yunt et al.	8,935,660 B2 1/2015 Aman et al.			
	7,478,381 B2 1/2009 Roberts et al.	8,970,598 B1 3/2015 Li et al.			
	7,483,984 B1 1/2009 Jonker et al.	9,008,962 B2 4/2015 Bandyopadhyay et al.			
	7,512,937 B2 3/2009 Chari et al.	9,021,428 B2 4/2015 Kumar et al.			
	7,548,237 B2 6/2009 David et al.	9,032,378 B2 5/2015 Jazdzewski et al.			
	7,565,640 B2 7/2009 Shukla et al.	9,118,598 B2 8/2015 Li et al.			
	7,574,692 B2 8/2009 Herscu	9,134,962 B1 9/2015 Stephens et al.			
	7,617,289 B2 11/2009 Srinivasan et al.	9,141,365 B1 9/2015 Zander et al.			
		9,298,479 B2 3/2016 Anson			
		9,338,609 B2 5/2016 Heydlauf			
		9,350,864 B1 5/2016 Brezinsky et al.			
		9,380,425 B2 6/2016 Meredith et al.			
		9,414,416 B2 8/2016 Adarapu et al.			
		9,460,109 B1 10/2016 Hagerman			
		9,584,318 B1 2/2017 Yang et al.			
		9,584,965 B2 2/2017 Good et al.			
		9,645,814 B1 5/2017 Roque et al.			
		9,823,906 B2 11/2017 Holzleitner et al.			
		9,864,405 B2 1/2018 Trachtenberg et al.			
		9,891,910 B1 2/2018 Cumming			
		9,910,640 B2* 3/2018 De Magalhaes G06F 8/35			
		10,063,999 B2 8/2018 Cherry et al.			
		10,068,116 B2 9/2018 Good et al.			
		10,109,230 B2 10/2018 Trachtenberg et al.			
		10,120,397 B1 11/2018 Zakhor et al.			
		10,296,303 B2 5/2019 Laethem et al.			
		10,341,818 B2 7/2019 Vutukuri et al.			
		10,352,707 B2 7/2019 Kordari et al.			
		10,382,897 B1 8/2019 Lanes et al.			
		10,417,355 B1 9/2019 Torgerson et al.			
		10,540,423 B2 1/2020 Osindero et al.			
		10,542,382 B2 1/2020 Good et al.			
		10,600,256 B2 3/2020 Ziegler et al.			
		10,674,320 B2 6/2020 Vutukuri et al.			
		10,735,907 B2 8/2020 Lanes et al.			
		10,789,716 B2 9/2020 Mikawa et al.			
		10,877,470 B2 12/2020 Burd et al.			
		10,908,666 B2 2/2021 Wang et al.			
		10,962,655 B2 3/2021 Pan			
		11,199,412 B2 12/2021 Kordari et al.			
		11,228,645 B2 1/2022 Mukhopadhyay et al.			
		11,265,673 B2 3/2022 Edge et al.			
		2001/0055425 A1 12/2001 Chiu			
		2002/0078255 A1 6/2002 Narayan			
		2003/0011601 A1 1/2003 Itoh et al.			
		2004/0143822 A1 7/2004 Jager et al.			
		2005/0131924 A1 6/2005 Jones			
		2005/0188080 A1 8/2005 Motsinger et al.			
		2005/0268280 A1 12/2005 Fildebrandt			
		2006/0010215 A1 1/2006 Clegg et al.			
		2006/0031359 A1 2/2006 Clegg et al.			

(56)

References Cited

U.S. PATENT DOCUMENTS

2006/0122972 A1 6/2006 Keohane et al.
 2006/0174233 A1 8/2006 Minadakis
 2006/0189326 A1 8/2006 Black et al.
 2006/0248112 A1 11/2006 Williams et al.
 2007/0250819 A1 10/2007 Fjeldstad et al.
 2007/0285241 A1 12/2007 Griebenow et al.
 2008/0048879 A1 2/2008 Lipman
 2008/0057936 A1 3/2008 Oosawa
 2008/0126426 A1 5/2008 Manas et al.
 2008/0127303 A1 5/2008 Wrighton et al.
 2008/0154691 A1 6/2008 Wellman et al.
 2008/0154712 A1 6/2008 Wellman
 2009/0007253 A1 1/2009 Chung et al.
 2009/0031281 A1 1/2009 Zhang et al.
 2009/0094688 A1 4/2009 Roy
 2010/0039247 A1 2/2010 Ziegler et al.
 2010/0083287 A1 4/2010 Maximilien et al.
 2010/0151943 A1 6/2010 Johnson et al.
 2010/0251155 A1 9/2010 Shah et al.
 2011/0072376 A1 3/2011 Moore et al.
 2011/0206023 A1 8/2011 Cohn et al.
 2012/0131660 A1 5/2012 Dalzell et al.
 2012/0260329 A1 10/2012 Suffling
 2012/0291006 A1 11/2012 Quine
 2013/0152041 A1 6/2013 Hatfield et al.
 2013/0160099 A1 6/2013 Fitzpatrick, III
 2013/0165143 A1 6/2013 Ziskind et al.
 2013/0219036 A1 8/2013 Gusev et al.
 2014/0033219 A1 1/2014 Gao et al.
 2014/0143252 A1 5/2014 Silverstein et al.
 2014/0195487 A1 7/2014 Behuria et al.
 2014/0207506 A1 7/2014 Palmert et al.
 2014/0207601 A1 7/2014 Soorianarayanan et al.
 2014/0280254 A1 9/2014 Feichtner et al.
 2014/0280493 A1 9/2014 Rodriguez et al.
 2014/0289700 A1 9/2014 Srinivasaraghavan et al.
 2015/0089615 A1 3/2015 Krawczyk et al.
 2015/0128104 A1 5/2015 Rath et al.
 2015/0193512 A1 7/2015 No et al.
 2015/0227406 A1 8/2015 Jan et al.
 2015/0286495 A1 10/2015 Lee
 2015/0301824 A1 10/2015 Patton et al.
 2015/0309770 A1 10/2015 Findlay et al.
 2015/0379257 A1 12/2015 Ahmed et al.
 2016/0026837 A1 1/2016 Good et al.
 2016/0094483 A1 3/2016 Johnston et al.
 2016/0132774 A1 5/2016 Prieditis
 2016/0142517 A1 5/2016 Chor
 2016/0295360 A1 10/2016 Jones et al.
 2016/0301769 A1 10/2016 Trost et al.
 2016/0381118 A1 12/2016 Andrews et al.
 2017/0006135 A1 1/2017 Siebel et al.
 2017/0034292 A1 2/2017 George
 2017/0118268 A1 4/2017 Wu et al.
 2017/0124366 A1 5/2017 Good et al.
 2017/0171024 A1 6/2017 Anerousis et al.
 2017/0223057 A1 8/2017 Amiri
 2017/0228447 A1 8/2017 Catania et al.
 2017/0243533 A1 8/2017 Trachtenberg et al.
 2017/0273864 A1 9/2017 Kaufman et al.
 2017/0289293 A1 10/2017 Rubtsov et al.
 2017/0308856 A1 10/2017 Bursey
 2017/0315787 A1* 11/2017 Waggoner G06F 8/73
 2017/0370728 A1 12/2017 Kordari et al.
 2018/0157384 A1 6/2018 Baneva et al.
 2018/0160271 A1 6/2018 Vutukuri et al.
 2018/0176097 A1 6/2018 Russell et al.
 2018/0284293 A1 10/2018 Pan
 2018/0321043 A1 11/2018 Kordari et al.
 2019/0042555 A1 2/2019 Jandhyala et al.
 2019/0082034 A1 3/2019 Hjerrild et al.
 2019/0082298 A1 3/2019 Good et al.
 2019/0132419 A1 5/2019 Wang et al.
 2019/0230474 A1 7/2019 Lanes et al.
 2019/0253843 A1 8/2019 Lanes et al.

2019/0349714 A1 11/2019 Vutukuri et al.
 2020/0004598 A1 1/2020 Brebner
 2020/0175778 A1 6/2020 Ziegler et al.
 2020/0228928 A1 7/2020 Good et al.
 2020/0285464 A1 9/2020 Brebner
 2020/0374391 A1 11/2020 Murgrabia et al.
 2021/0293977 A1 9/2021 Pan
 2022/0044055 A1 2/2022 Jean et al.

OTHER PUBLICATIONS

Kim, S. et al., "A Configuration Management System for Multi-Hop Zigbee Networks," IEEE, 2009, 2 pages.
 Liu, L. et al., "Cloud Service Portal for Mobile Device Management," IEEE International Conference on E-Business Engineering, 2010, pp. 474-478.
 Petrucci, V. et al., "Optimized Management of Power and Performance for Virtualized Heterogenous Server Clusters," 2011 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2011, pp. 23-32.
 Sinz, C. et al., "Verifying CIM Models of Apache Web-Server Configurations," Proceedings of the Third International Conference on Quality Software (QSIC '03), IEEE, 2003, 8 pages.
 Totok, A. et al., "Exploiting Service Usage Information for Optimizing Server Resource Management," ACM Transactions on Internet Technology, vol. 11, No. 1, Article 1, Jul. 2011, pp. 1-26.
 Brown, D.K. et al., "JMS: An Open Source Workflow Management System and Web-Based Cluster Front-End for High Performance Computing," PLOS One, Aug. 17, 2015, pp. 1-25.
 Li, Y. et al., "Consistency Aware Update Schedule in Multi-Server Distributed Virtual Environments," ACM, 2010, pp. 1-8.
 Pan, H. et al., "Design and Implementation of Server Management System Based on Docker," 2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA), 2019, pp. 48-52.
 Park, J.G. et al., "Cluster Management in a Virtualized Server Environment," ACM, 2008, pp. 2211-2214.
 Chan, M. et al., "Extracting Web Design Knowledge: The Web De-Compiler," IEEE, 1999, pp. 547-552.
 Ehrig, K. et al., "Generation of Visual Editors as Eclipse Plug-Ins," ACM, 2005, pp. 134-143.
 Gallego-Carrillo, M. et al., "SOTA: a Visualization Tool for Symbol Tables," ACM, 2005, p. 385.
 Gotz, D. et al., "Interactive Visual Synthesis of Analytic Knowledge," IEEE Symposium on Visual Analytics Science and Technology, 2006, pp. 51-58.
 International Search Report and Written Opinion dated Oct. 30, 2018 for PCT International Application No. PCT/US2018/035953, 16 pages.
 Kannan et al., "Digital Library for Mulsemedia Content Management," ACM, pp. 275-276 (Year: 2010).
 Key, A. et al., "VizDeck: Self-Organizing Dashboards for Visual Analytics," ACM, 2012, pp. 681-684.
 Luo et al., "Application of Digital Content Management System Based on Data Warehouse and Data Mining Technology," IEEE, pp. 988-992 (Year: 2012).
 Lyu, M.R. et al., "A Wireless Handheld Multi-modal Digital Video Library Client System," ACM, 2003, pp. 231-238.
 Mancini, R., "Interacting with a Visual Editor," ACM, 1996, pp. 125-131.
 Matsubara et al., "Managing a Media Server Content Directory in Absence of Reliable Metadata," IEEE, pp. 873-877 (Year: 2009).
 Meibner, M. et al., "Enabling Classification and Shading for 3D Texture Mapping Based vol. Rendering Using OpenGL and Extensions," ACM, 1999, retrieved online on Nov. 24, 2019, pp. 201-214 and 256.
 Nam, B. et al., "Mobile Based Digital Contents Conversion Service Implementation," ACM, 2009, pp. 1432-1436.
 PCT International Search Report and Written Opinion dated Jun. 12, 2019 for International Application No. PCT/US2018/063849, 21 pages.
 Pop et al., "On Cloud Development for Digital Preservation Environments," IEEE, pp. 443-444 (Year: 2014).

(56)

References Cited

OTHER PUBLICATIONS

Raschke, M. et al., "Annotation of Graphical Elements in Visualizations for an Efficient Analysis of Visual Tasks," ACM, 2014, pp. 1387-1392.

Stanchev, P. et al., "Enhanced User Experience and Behavioral Patterns for Digital Cultural Ecosystems," ACM, 2017, retrieved online on Nov. 8, 2018, pp. 287-292, <http://delivery.acm.org/10.1145/3170000/3167063/p287-stanchev.pdf>.

Vollebregt, T. et al., "Declarative Specification of Template-Based Textual Editors," ACM, 2012, retrieved online on Nov. 9, 2019, 7 pages, <http://delivery.acm.org/10.1145/2430000/2427056/a8-vollebregt.pdf>.

Extended European Search Report dated Feb. 9, 2021 for EP Application No. 18812640.3, 6 pages.

Brenna, L. et al., "Configuring Push-Based Web Services," IEEE, 2005, 6 pages.

Fitzgerald, S. et al., "A Directory Service for Configuring High-Performance Distributed Computations," IEEE, 1997, pp. 365-375.

Yalagandula, P. et al., "A Scalable Distributed Information Management System," ACM, 2004, pp. 379-390.

Hou, K. et al., "Indoor Location Based on WiFi," Proceedings of the 2015 IEEE 19th International Conference on Computer Supported Cooperative Work in Design (CSCWD), 2015, pp. 331-336.

Haque, I.T. et al., "Profiling-Based Indoor Localization Schemes," IEEE Systems Journal, vol. 9, No. 1, Mar. 2015, pp. 76-85.

Xiao, T.T. et al., "Study Of Fingerprint Location Algorithm Based On WiFi Technology For Indoor Localization," IEEE, Sep. 2014, pp. 604-608.

Yun, S. et al., "Data-centric Middleware based Digital Twin Platform for Dependable Cyber-Physical Systems," IEEE, Jul. 2017, pp. 922-926.

Autiosalo, J. et al., "Twinbase: Open-Source Server Software for the Digital Twin Web," IEEE Access, vol. 9, Oct. 2021, pp. 140779-140798.

Yassin, M. et al., "A Survey of Positioning Techniques and Location Based Services in Wireless Networks," IEEE, 2015, 5 pages.

U.S. Appl. No. 16/566,196, filed Sep. 10, 2019, Brebner.

U.S. Appl. No. 16/566,281, filed Sep. 10, 2019, Brebner.

U.S. Appl. No. 16/818,269, filed Mar. 13, 2020, Brebner.

U.S. Appl. No. 16/818,326, filed Mar. 13, 2020, Brebner.

U.S. Appl. No. 16/818,370, filed Mar. 13, 2020, Brebner.

U.S. Appl. No. 16/818,411, filed Mar. 13, 2020, Brebner.

Cullen, G. et al., "Cooperatively Extending the Range of Indoor Localisation," ISSC, LYIT Letterkenny, Jun. 20-21, 2013, 8 pages.

Whiting, S. et al., "Geolocation From Received Signal Strength," 2018 52nd Asilomar Conference on Signals, Systems, and Computers, IEEE, Oct. 2018, pp. 1147-1151.

* cited by examiner

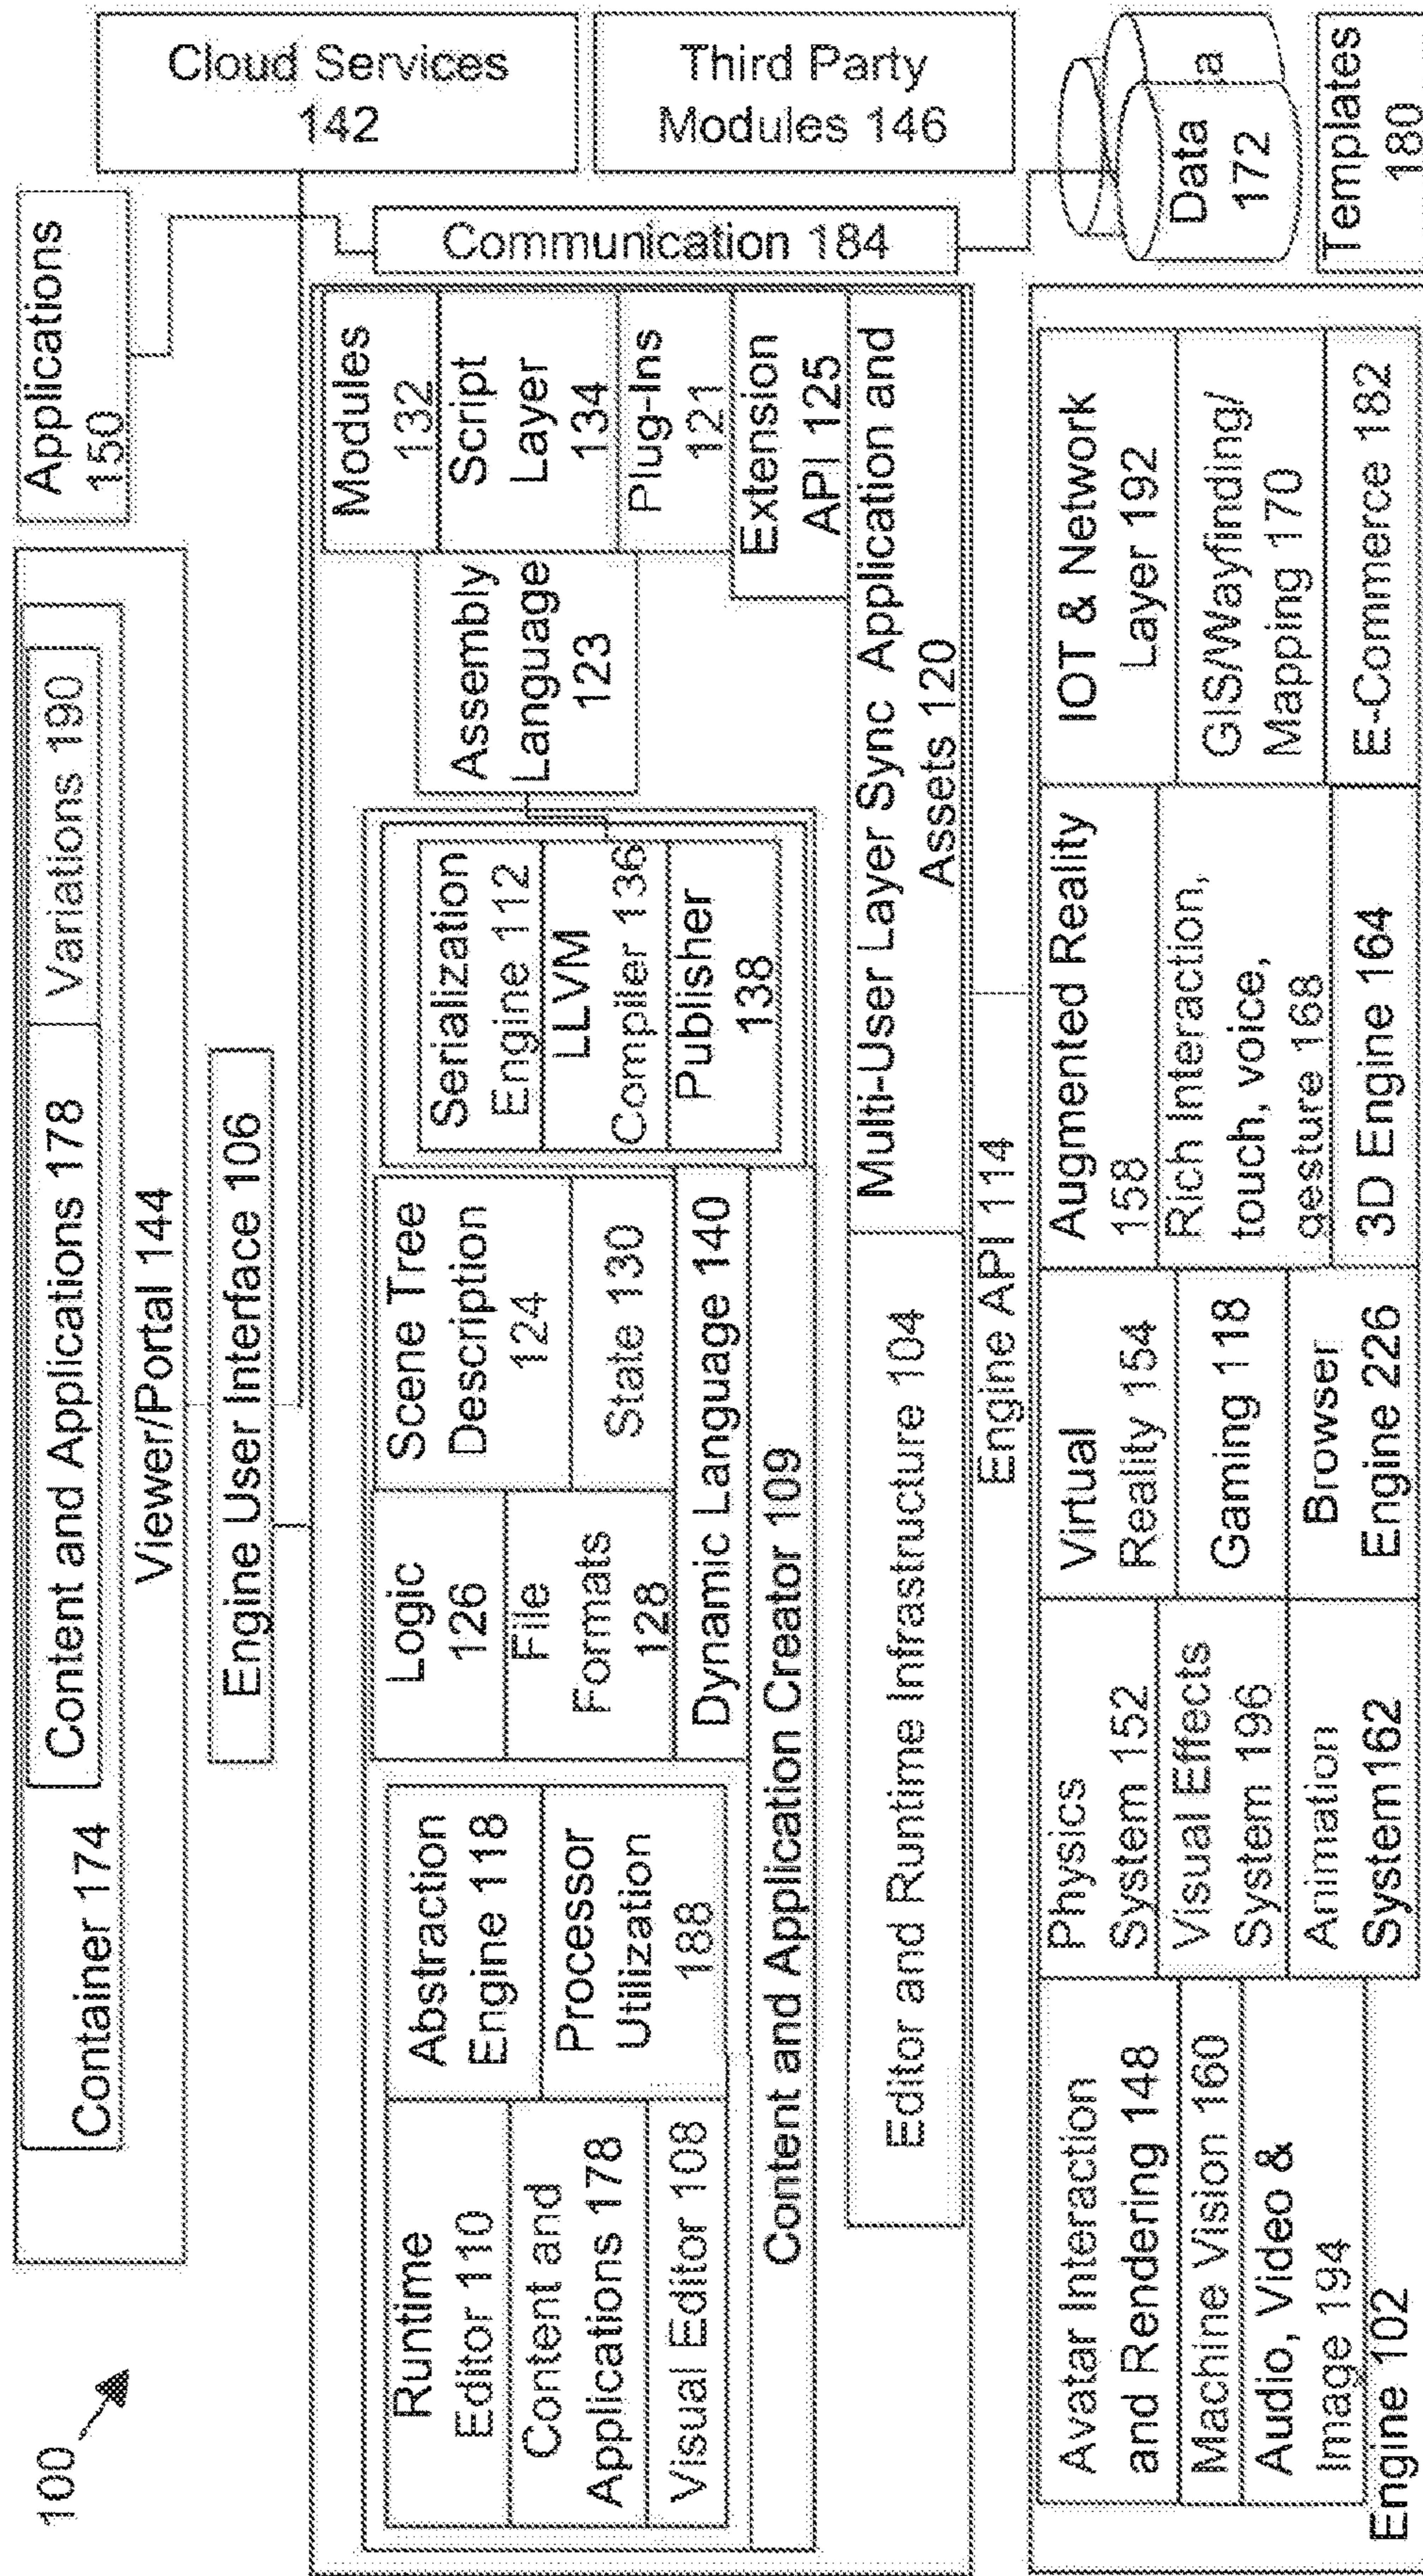


FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A




FIG. 2B




FIG. 2C

300

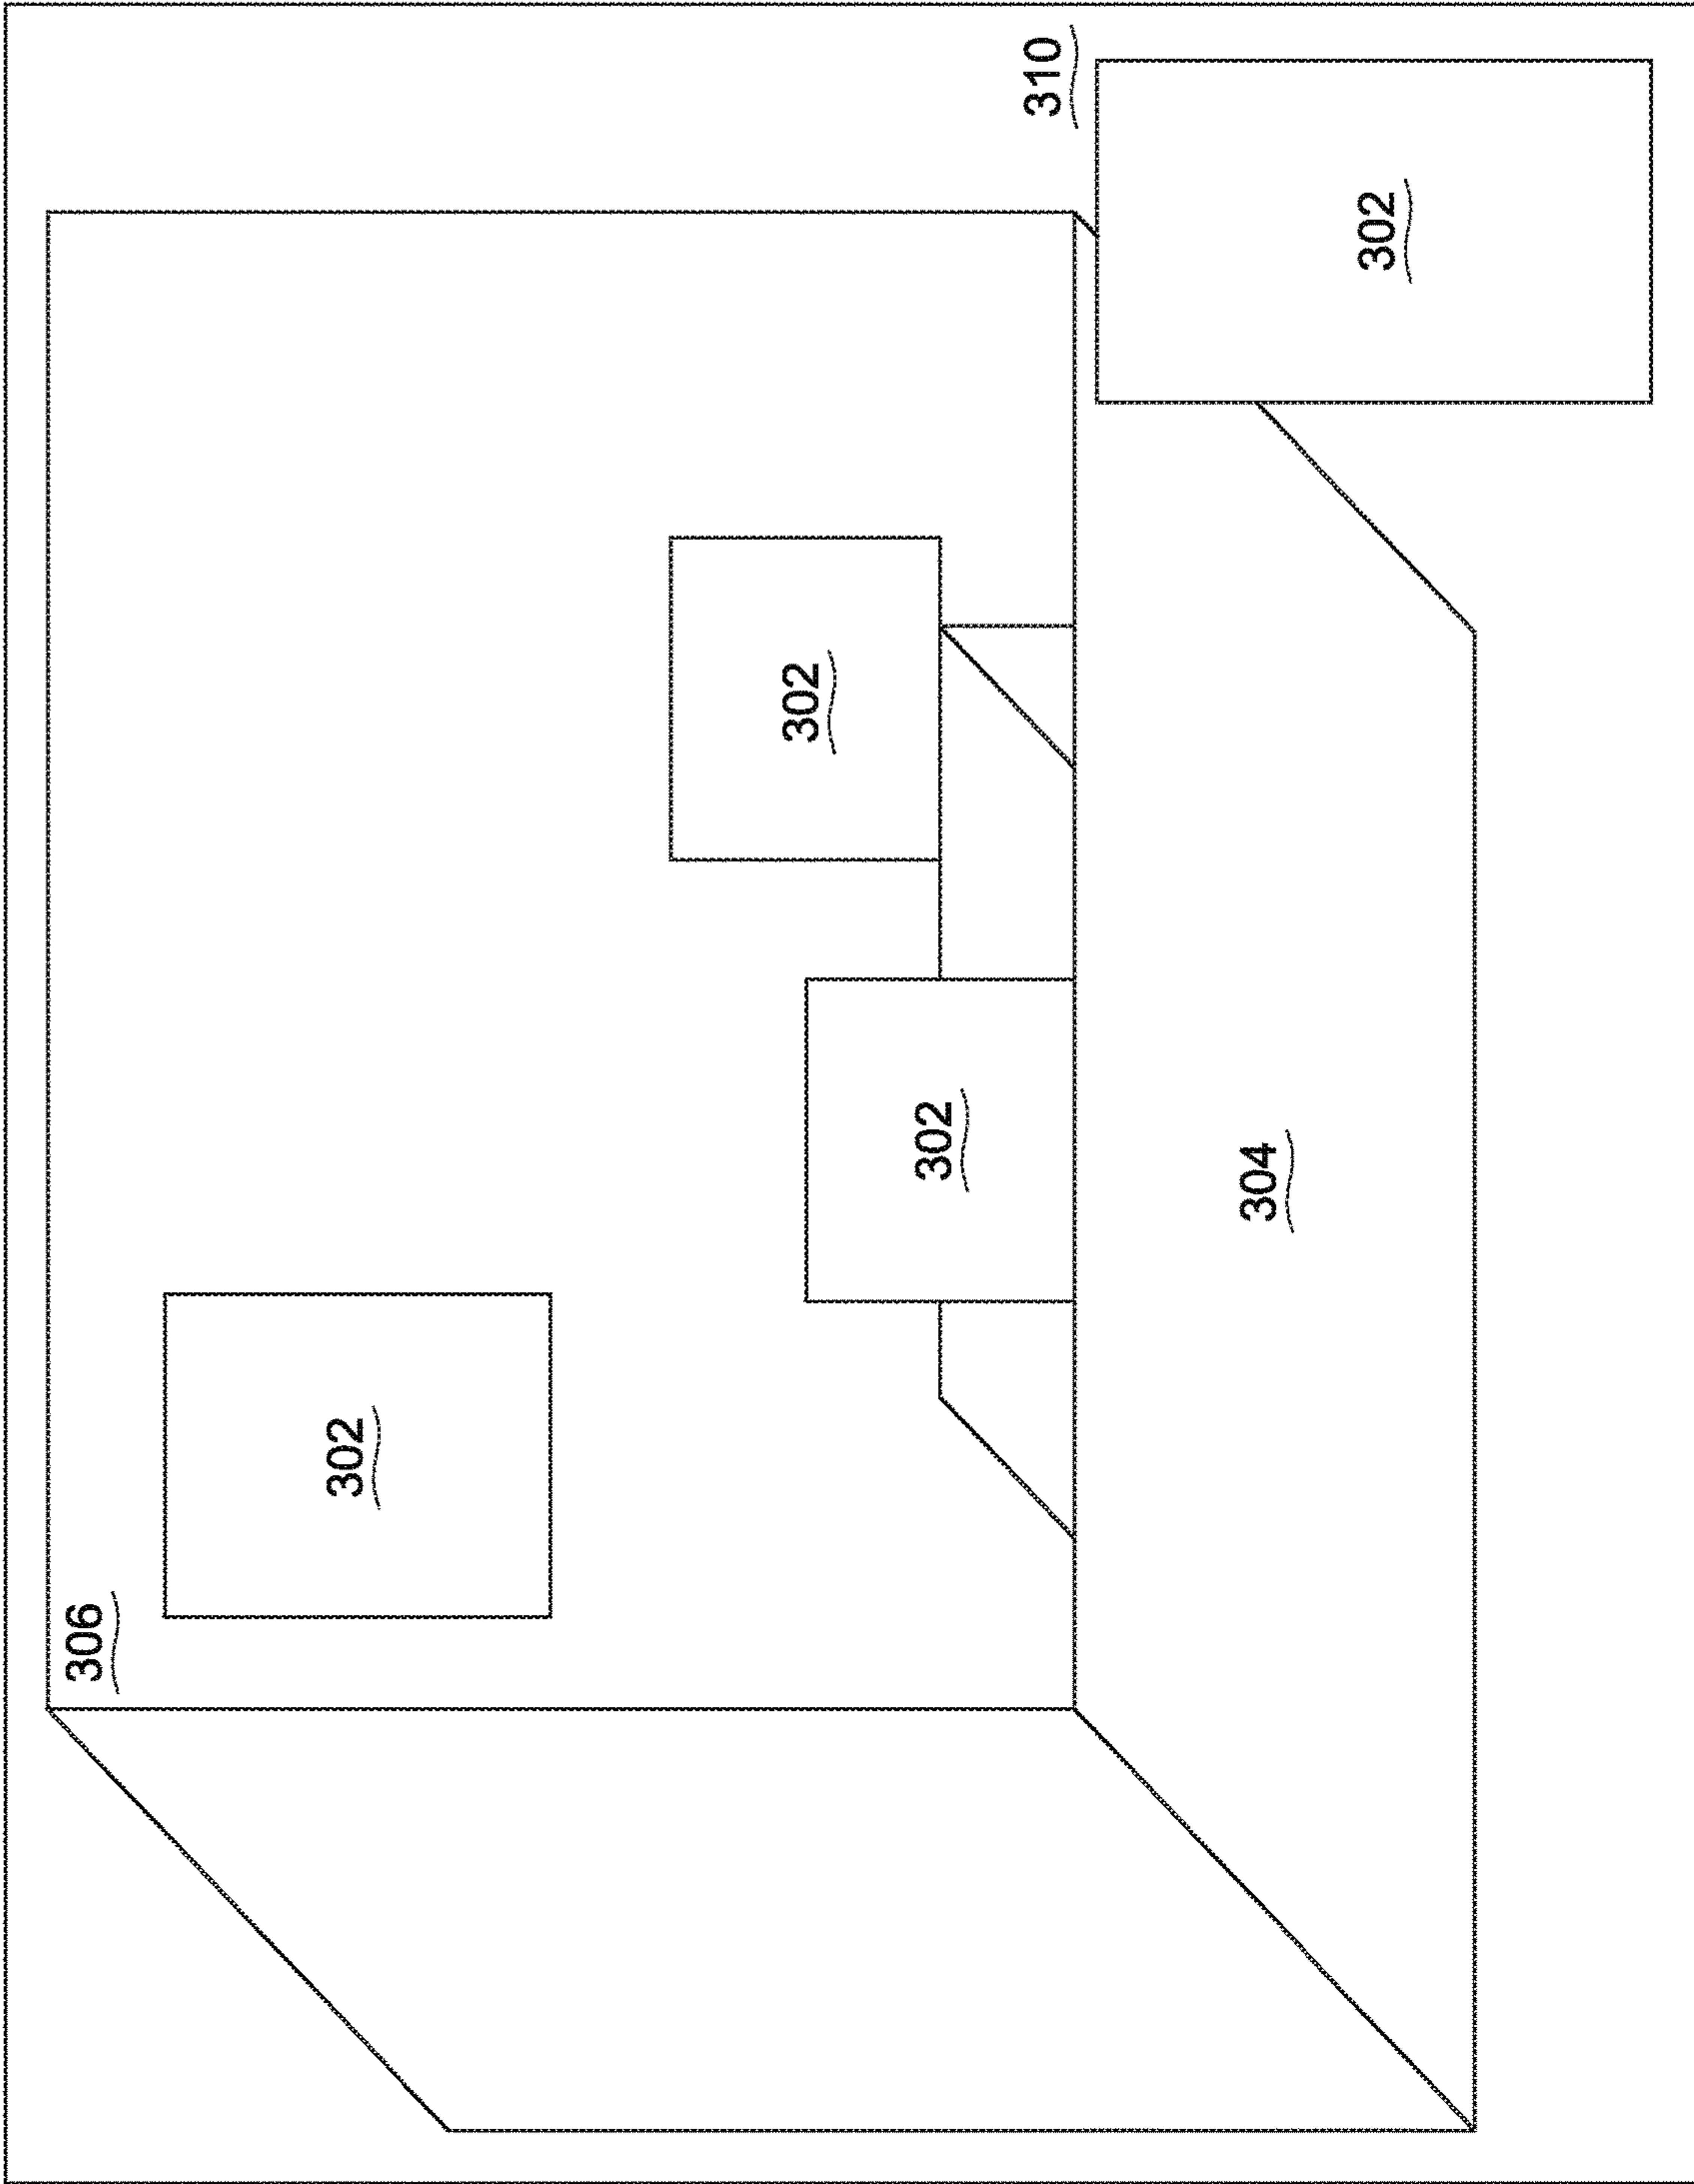


FIG. 2D




FIG. 2E




FIG. 2F

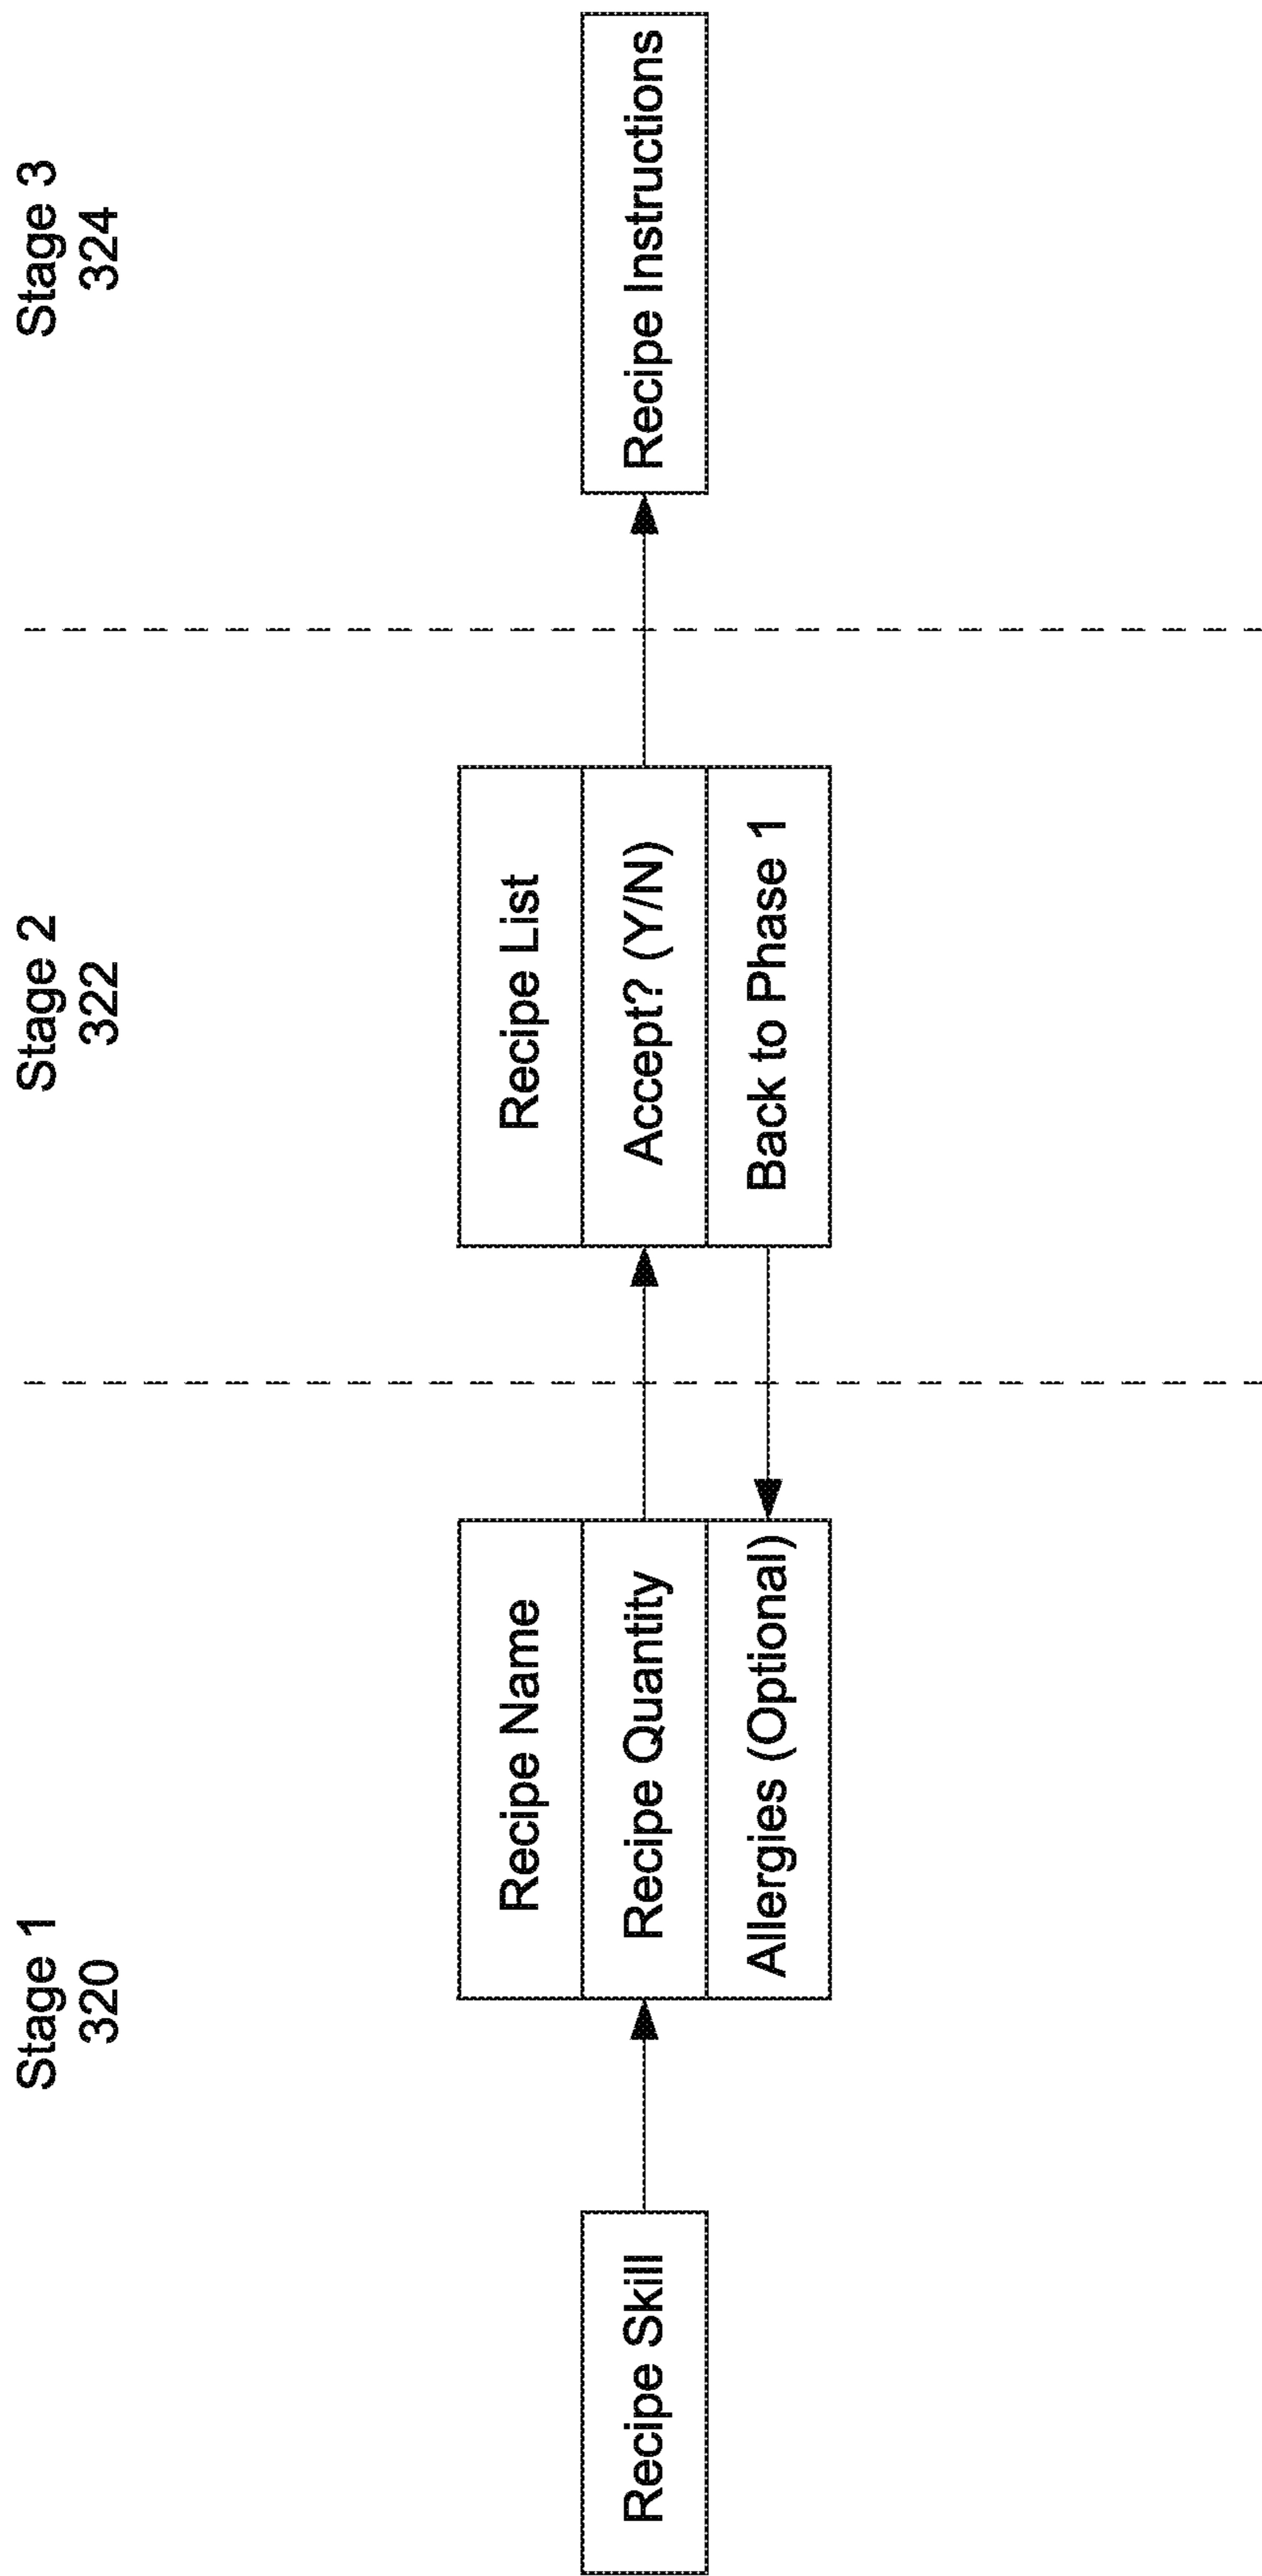


FIG. 2G




FIG. 3

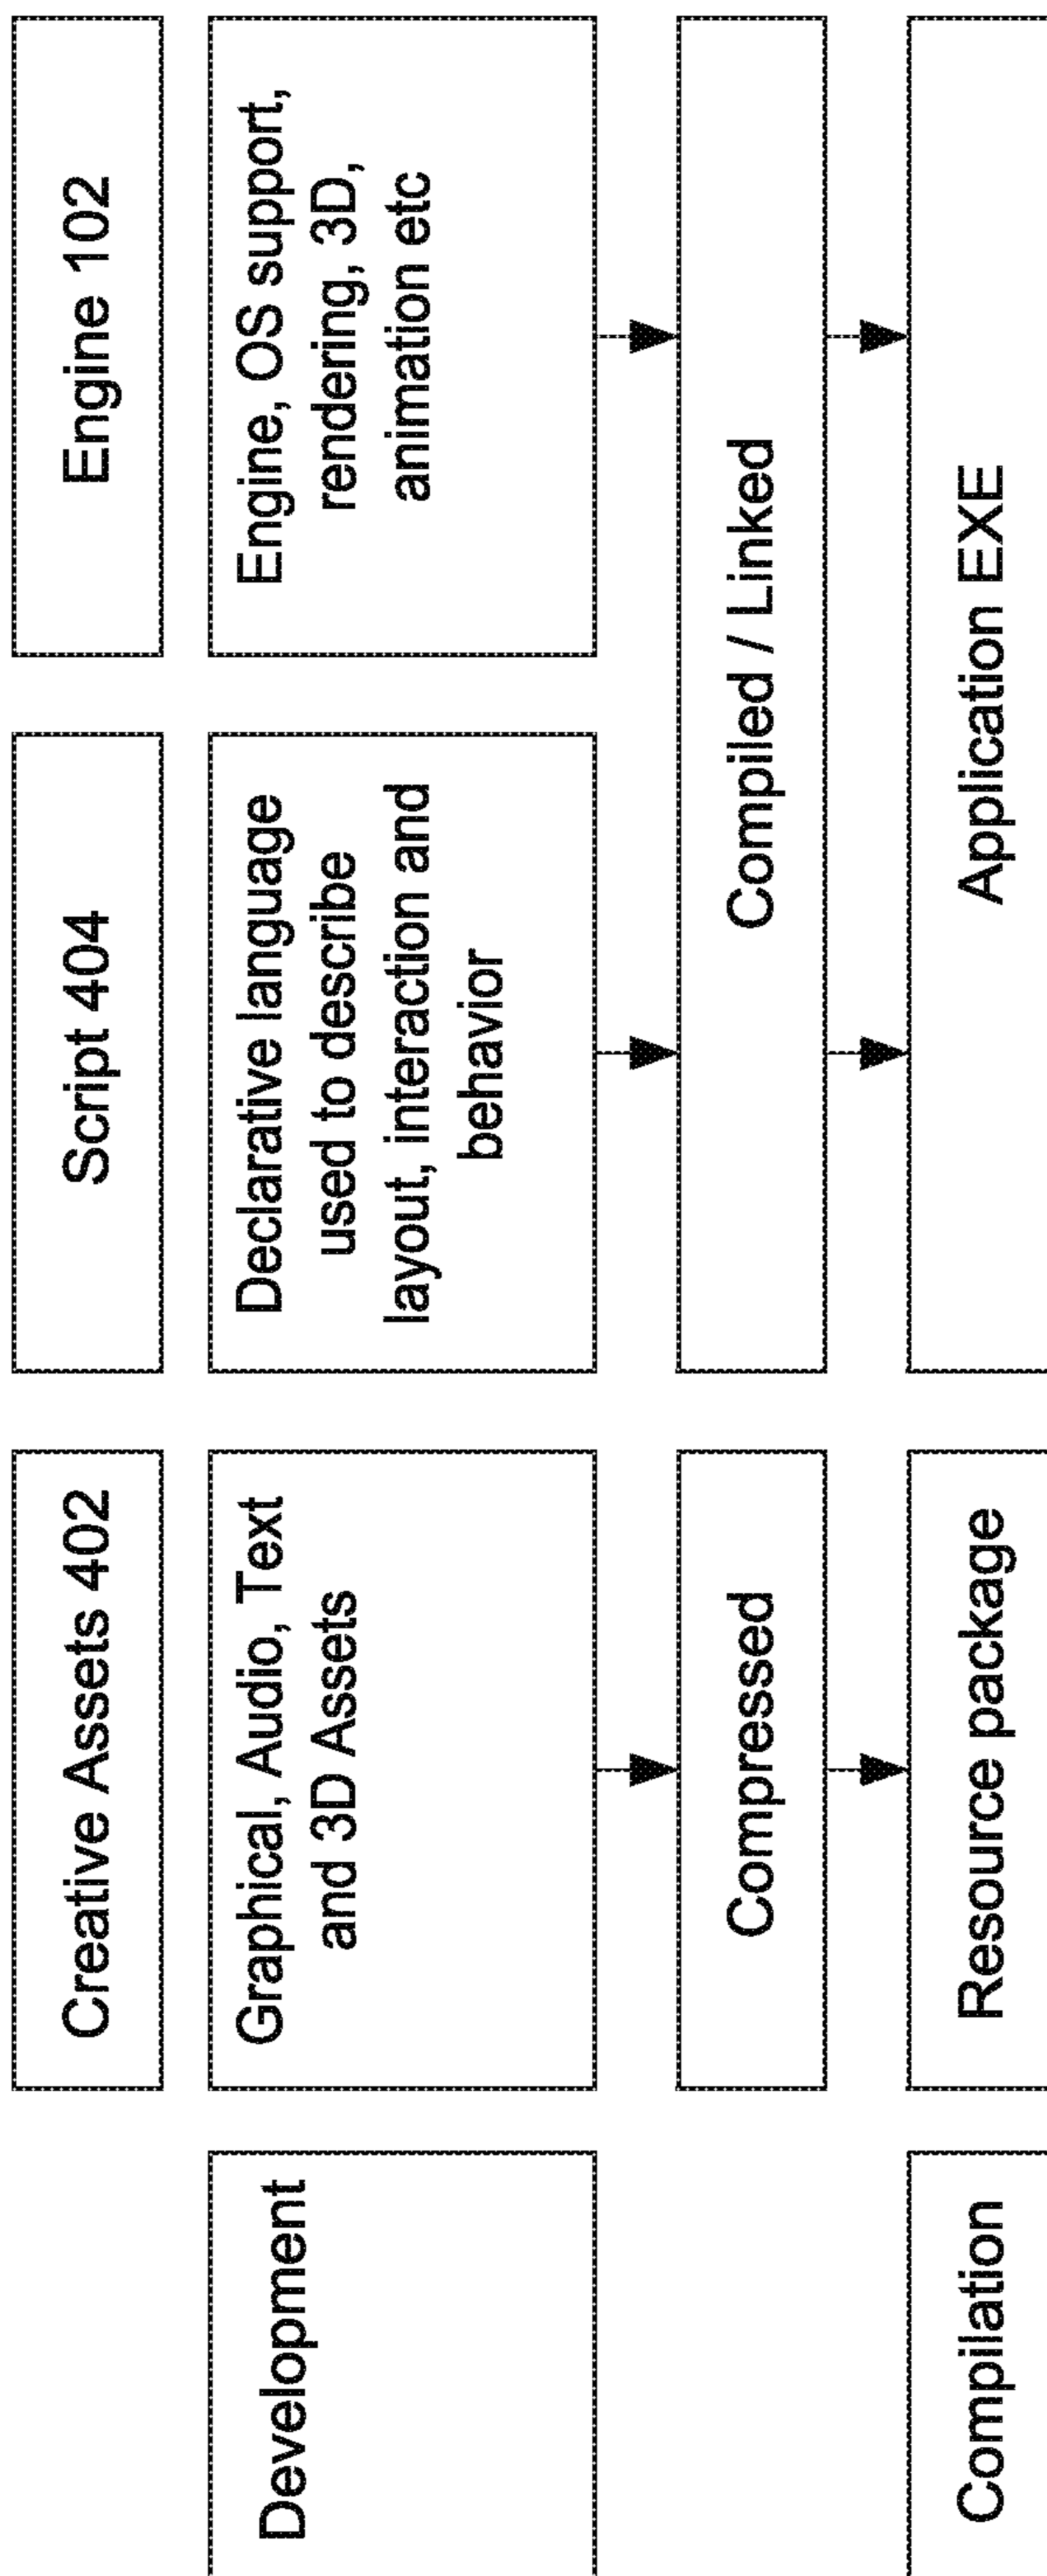


FIG. 4


```
instance image apple
    .filename='apple.png'
end
```

FIG. 5

```
instance image apple
    .filename='apple.png'
instance image worm
    .filename'worm.png'
end
end
```

FIG. 6

```
define image button
    .filename='button.png'
end
```

FIG. 7A

```
define image button
    .filename='button.png'
    event on_release()
    method on_mouse_down(int x, int y, int mod)
        raise .on_release()
    end
end
instance button mybutton
    method on_release()
    end
end
```

FIG. 7B

```
Real velocity = 2.3      (2.3 is the expression, the = is the  
assignment)  
Real speed = velocity * 1.87      (velocity * 1.87 is the  
expression)  
Int length = sqrt((x1-x2)^2+(y1-y2)^2) (sqrt is a method inside  
the expression)
```

FIG. 8

```
if velocity>2.1  
    //place code here  
end
```

FIG. 9A

```
if velocity>2.1
  //place code here
elseif velocity==1.5
  //place code here
else
  //place code here
end
```

FIG. 9B

```
int a=0
loop a<100
  //do something many times...
  a++
end
```

FIG. 10A

```
int a=0
loop
  //do something many times...
  a++
end a>100
```

FIG. 10B

```
loop int i=0, i<5, i++
  print(i)
end
```

FIG. 10C

```
map m
m['b']='beta'
m[0]='zero'
m['a']='alpha'
m[1]='one'
//loop forward
loop var k=m.first_key(), k.is_valid(), k=m.next_key(k)
  print (k->+m[k])
end
```

FIG. 10D

```
string my_big_array[][]  
my_array[1][ "name" ] = "cow"  
my_array[1][ "age" ] = "13"  
my_array[2][ "name" ] = "sheep"  
my_array[2][ "age" ] = "3"
```

FIG. 11A

```
int my_array[]  
int a=0  
loop a<100  
    my_array[a]=rnd(100)  
    a++  
end
```

FIG. 11B

```
int my_array[]
```

FIG. 11C

```
string my_array[]  
my_array[1] = "cow"  
my_array[2] = "sheep"
```

FIG. 11D

```
instance button add_button
    .x=20
    .y=100
    .width=200
    .height=50
    .caption="add"

method on_mouse_down(int x, int y, int mod)
    do_add()
end

state language_german
    .width=300
    .x=5
    .caption="hinzufugen"
method on_mouse_down(int x, int y, int mod)
    check_region()
    do_add()
end

state orientation_landscape
    .width=400
end

end
```

FIG. 12A

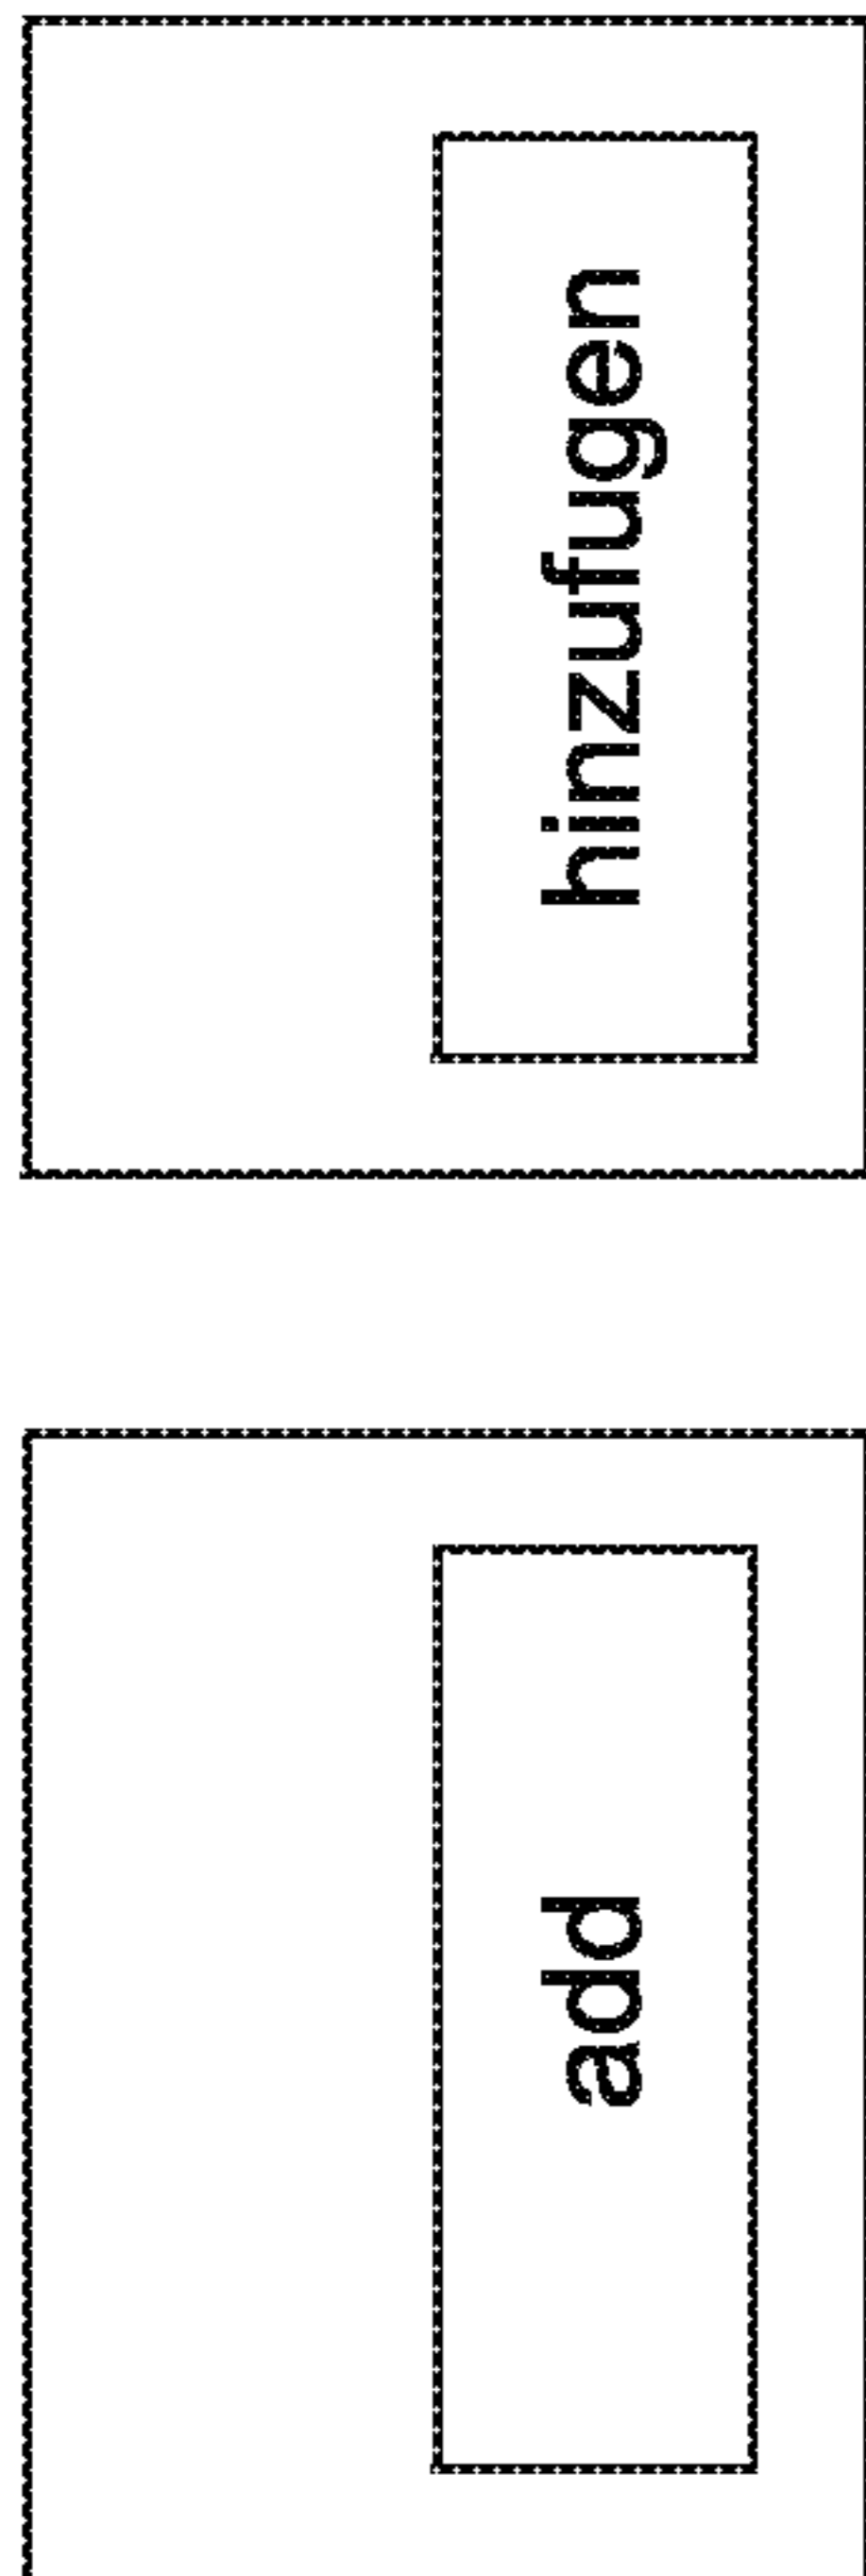


FIG. 12B

Rule	Target	Variation
OS=iOS	Page1: Object: Toolbar	Hide
Region=NZ	Page2: Object: BuyButton	tax=15.0
Region=NZ	Page2: Object: Buybutton	Custom Actions
Orientation=Portrait	Masterpage1: Region: column4	Hide
Orientation=Landscape	Masterpage: Region: toolbar	Height=10mm

FIG. 13

```
instance form myform
instance button mybutton
    .caption='scale and rotate'
    method on_release()
        myform.scale=0.5 in 1000 tween linear
        myform.rotate=20 in 1000 tween linear
end
end
end
```

FIG. 14

```
instance image button
    .caption='OK'
    .width=100
end
....
if language = 'german'
    button.caption='Ordnung'
    button.width=200
end
```

FIG. 15A

```
instance image button
    .caption='OK'
    .width=100
    stat lang_German
        .caption='Ordnung'
        .width=200
    end
end
```

FIG. 15B

```
instance image tree
instance image apple1
    .x=0
    .y=0
end
instance image apple2
    .x=100
    .y=200
end
end
```

FIG. 16

```
instance image tree
method on_click()
    .apple1.state=picked
end
instance image apple1
    .x=0
    .y=0
    .filename="apple.png"
state picked
    .x=100
    .filename="applecore.png"
end
end
end
```

FIG. 17

```
instance image tree
method on_click()
    .apple.x=100
end
instance image apple1
    .x=0
    .y=0
end
end
```

FIG. 18




FIG. 19




FIG. 20




FIG. 21




FIG. 22




FIG. 23A




FIG. 23B

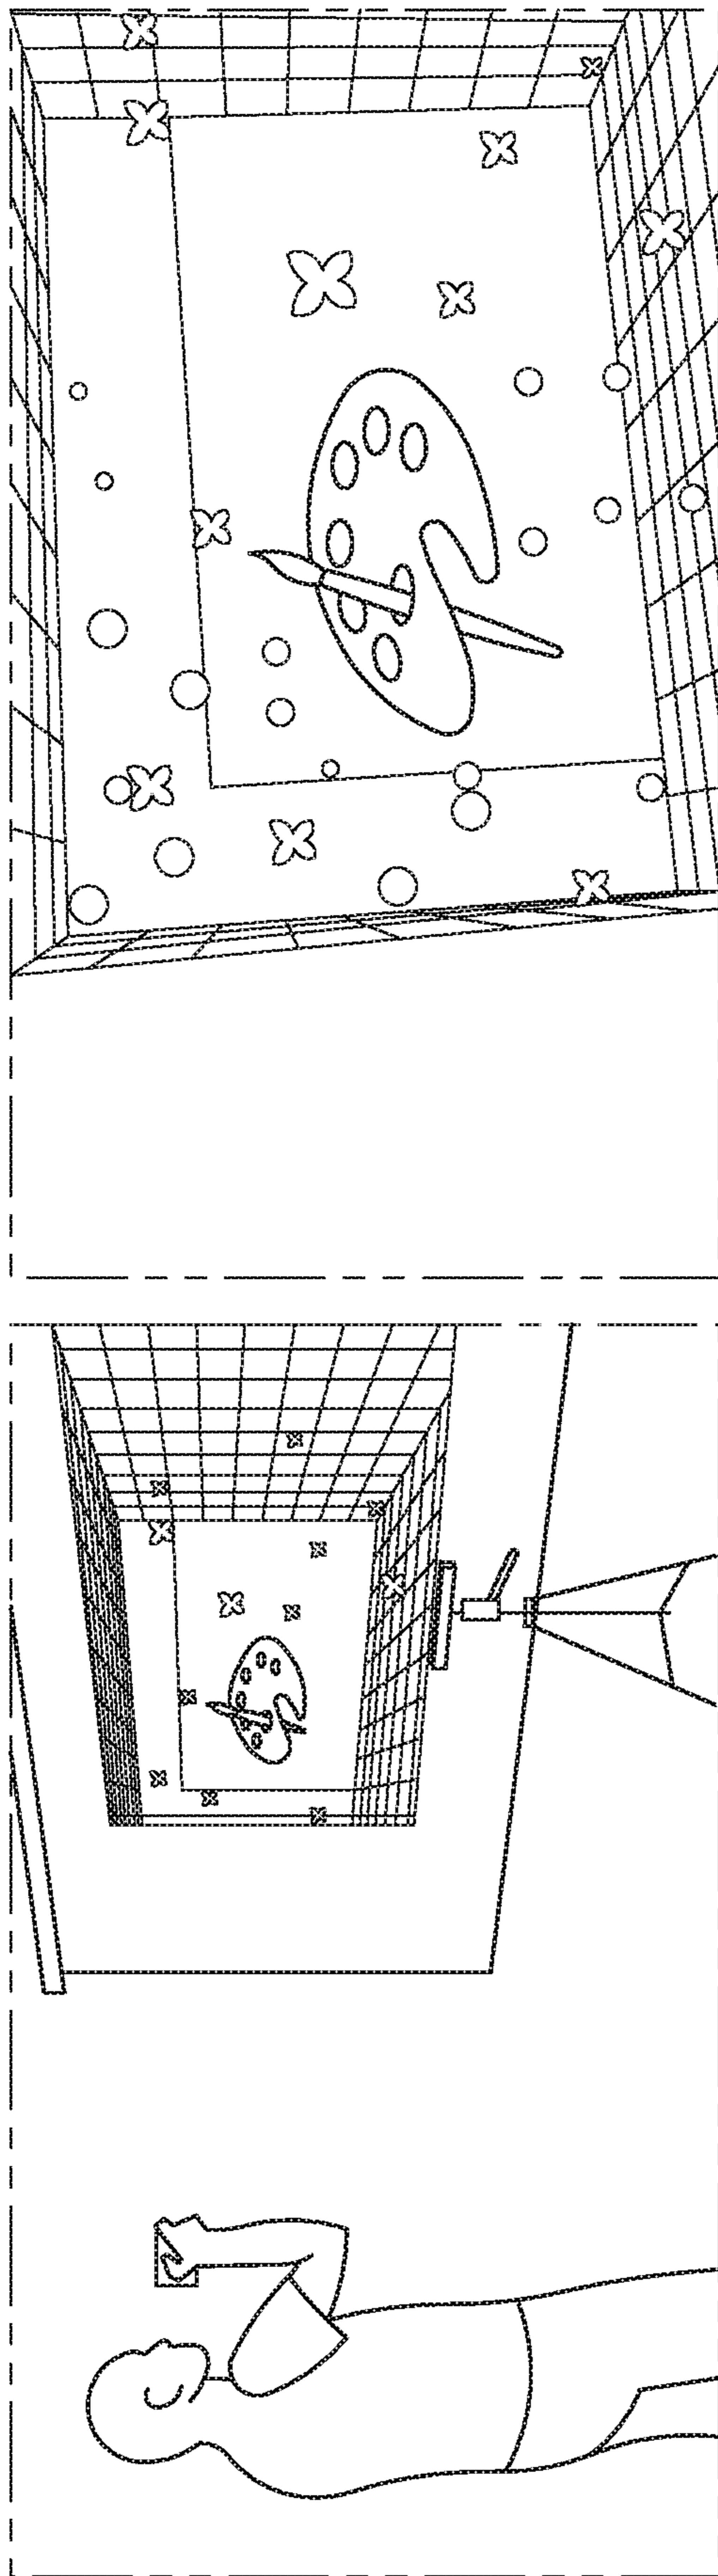


FIG. 24




FIG. 25




FIG. 26A




FIG. 26B




FIG. 27A




FIG. 27B




FIG. 28A




FIG. 28B




FIG. 28C




FIG. 28D




FIG. 29A




FIG. 29B

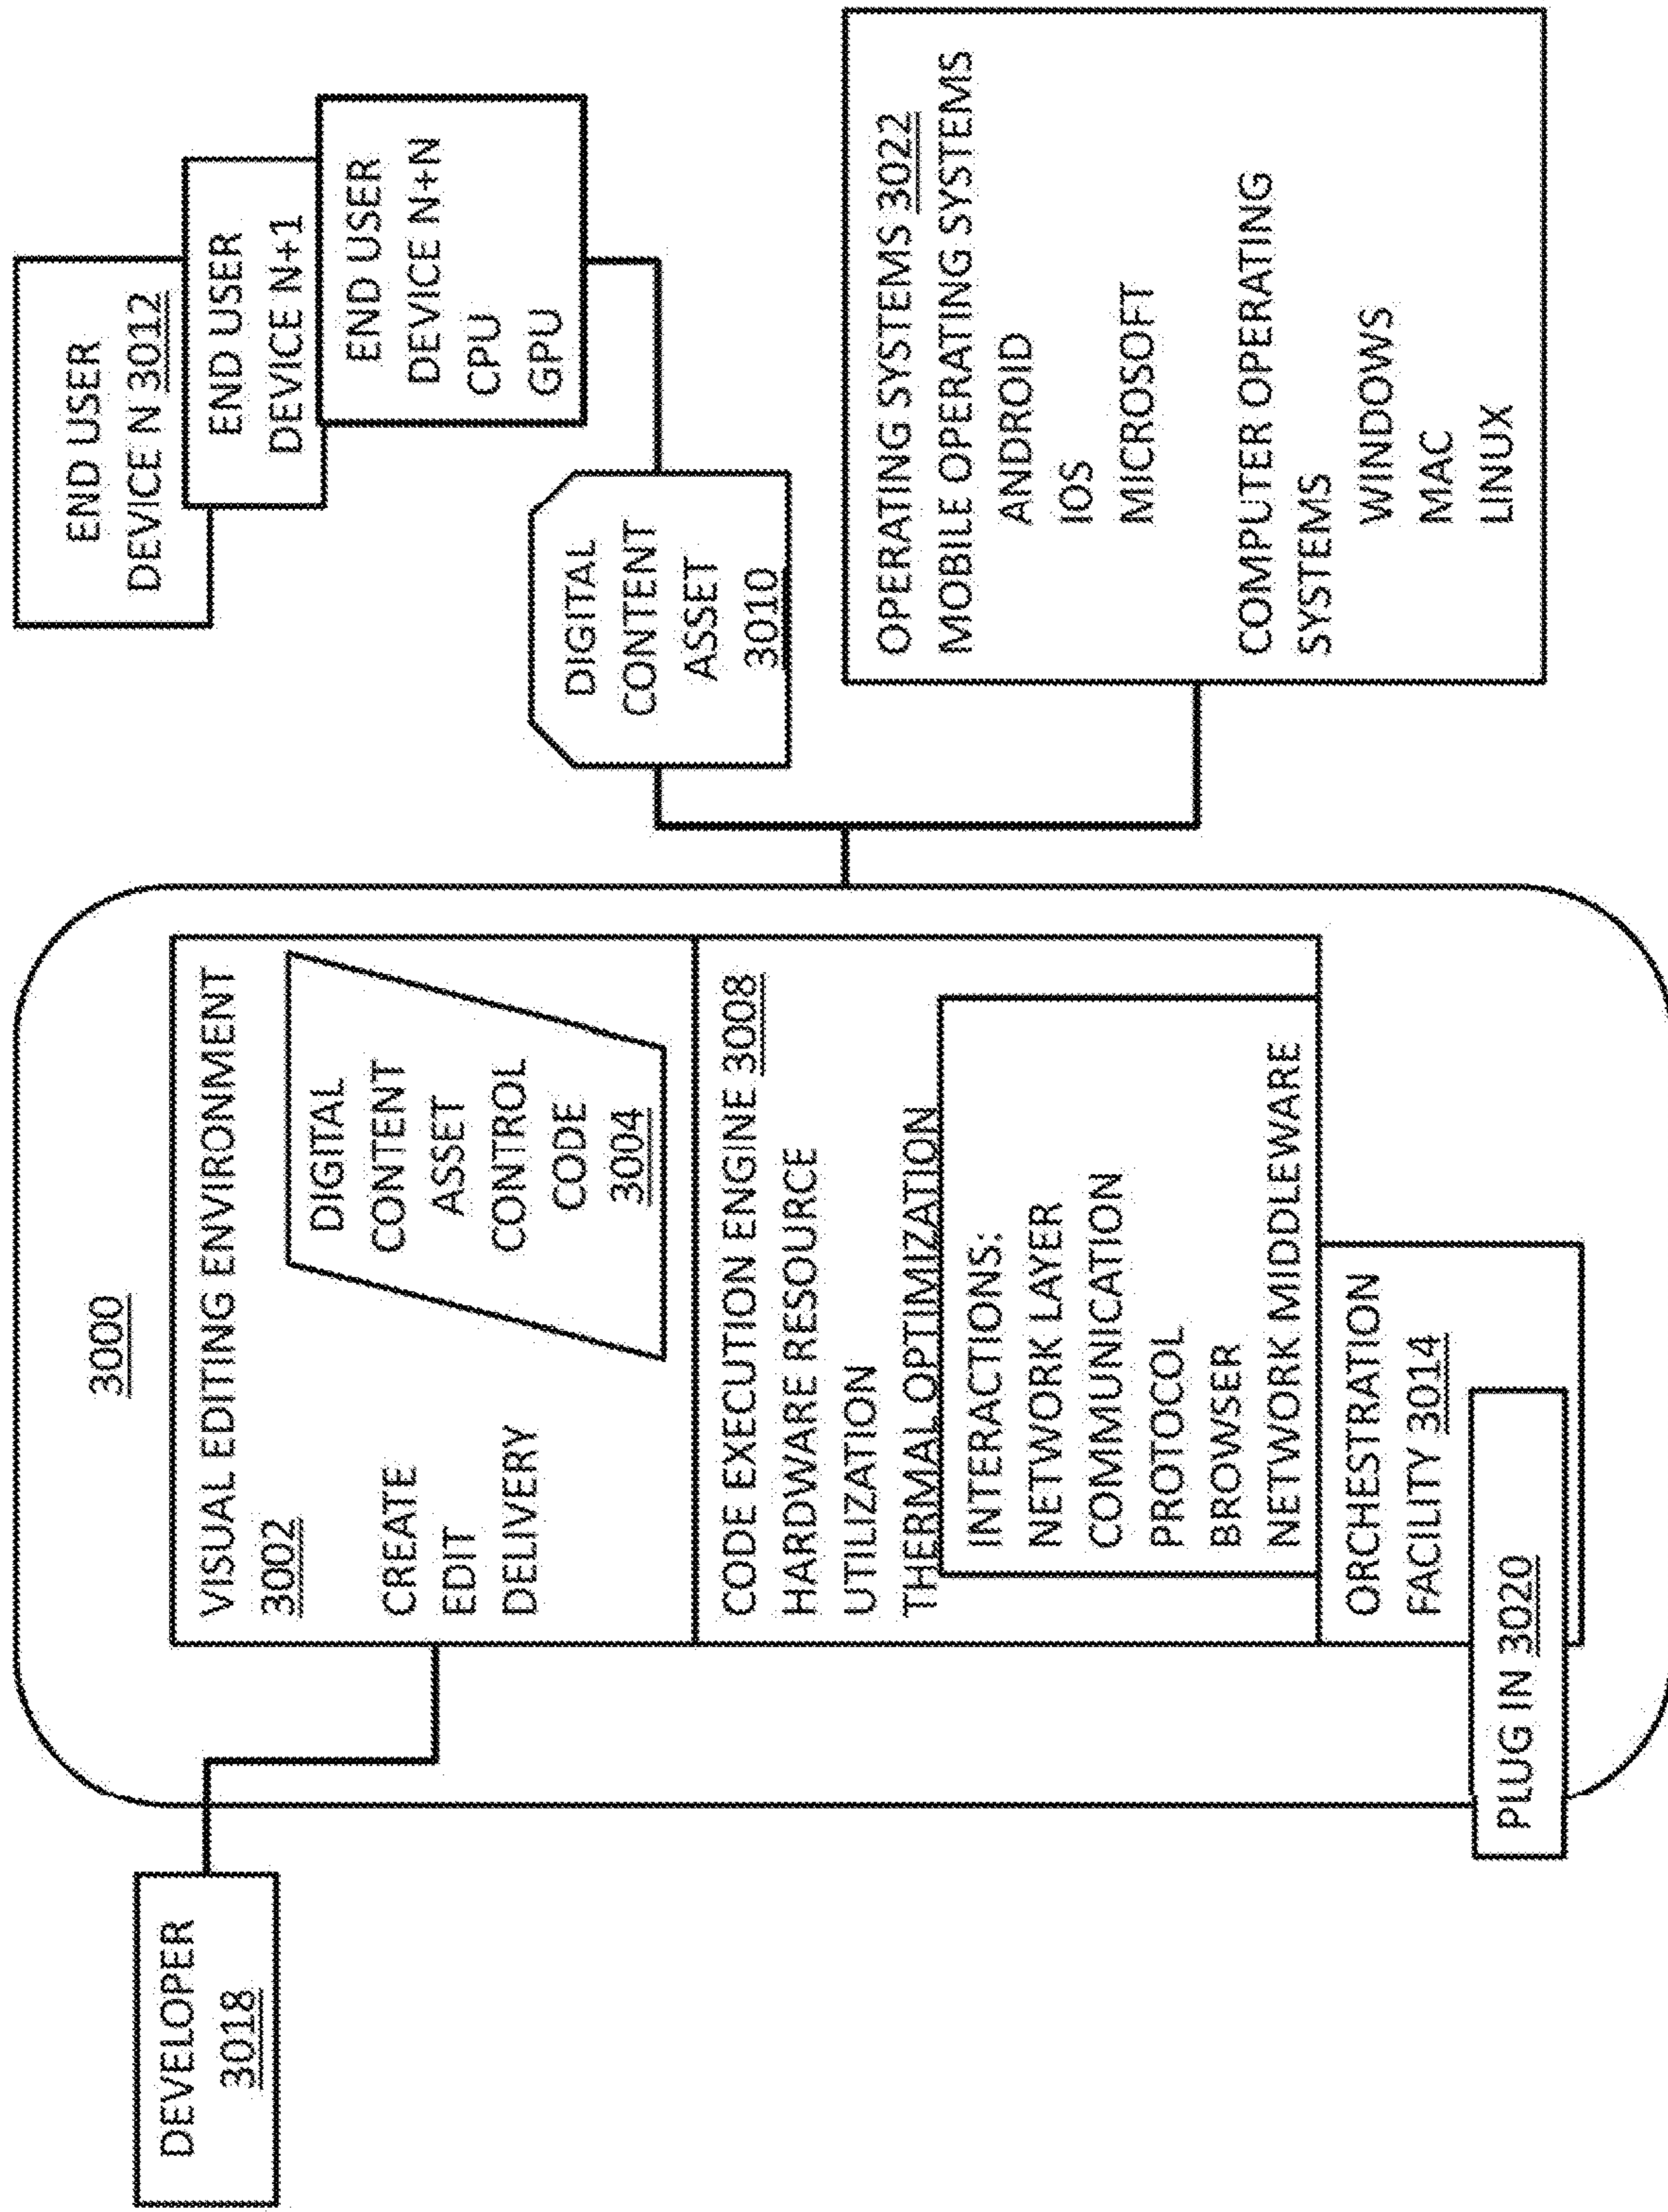


FIG. 30




FIG. 31

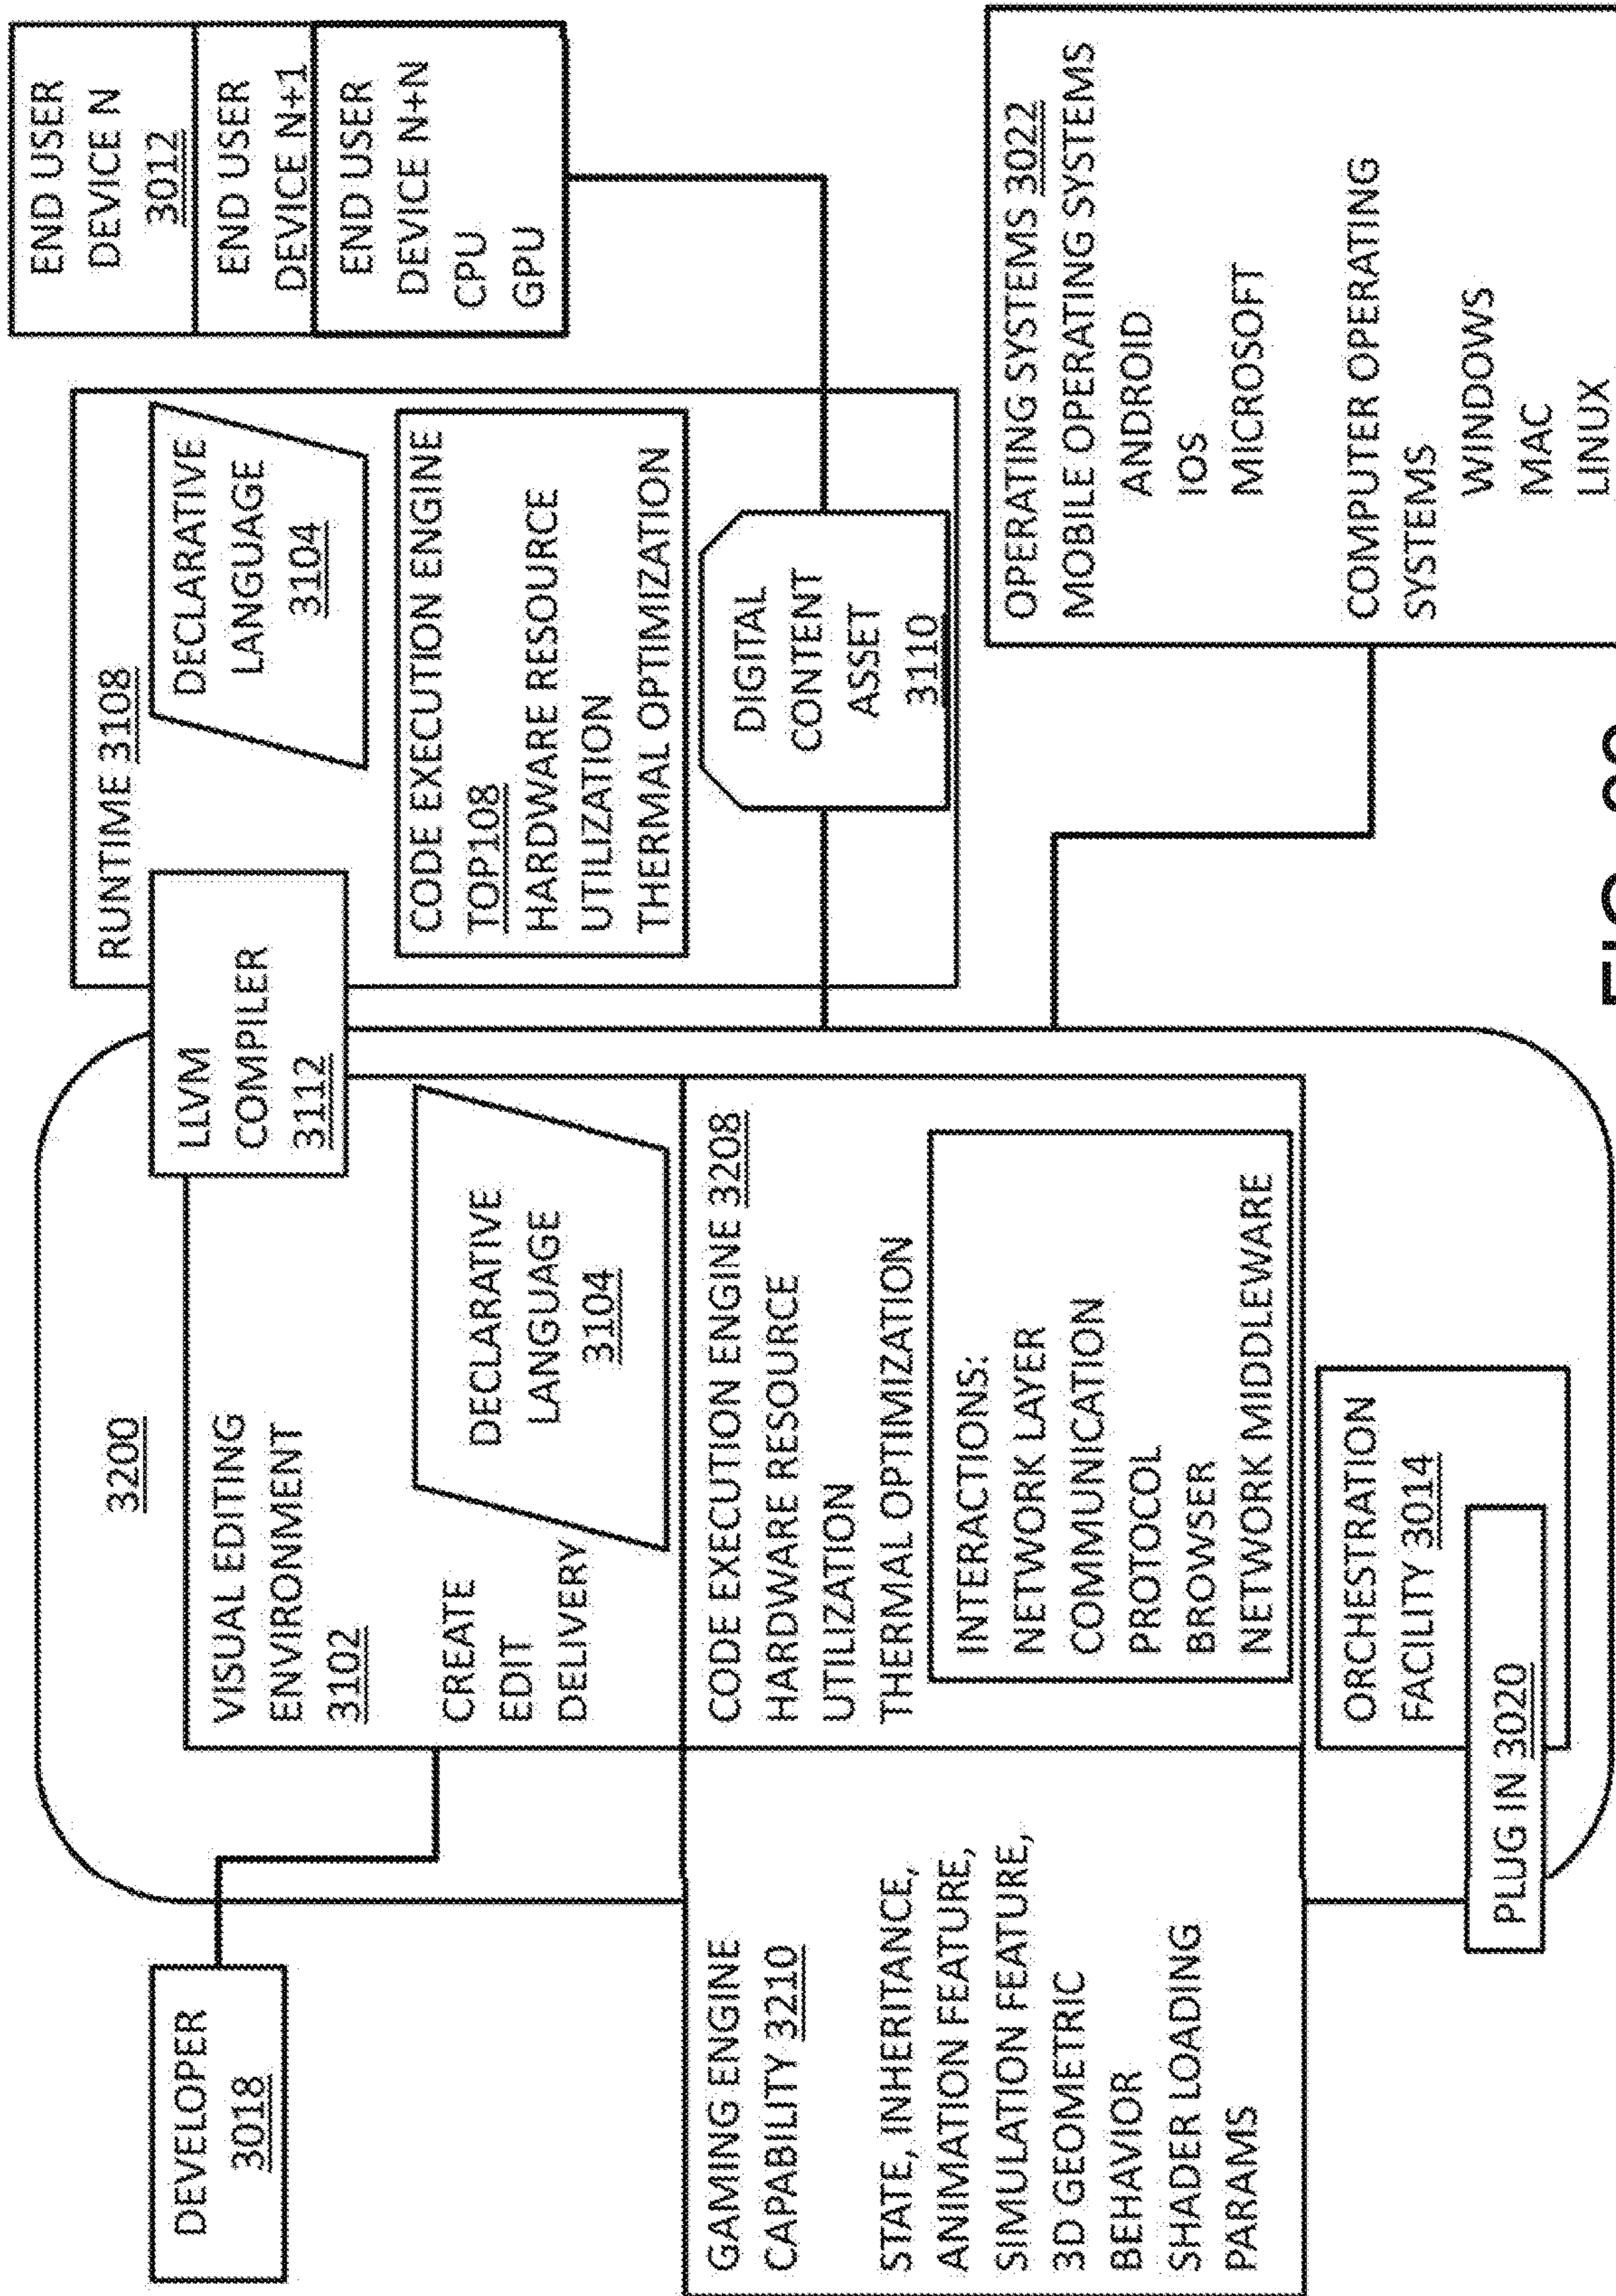


FIG. 32

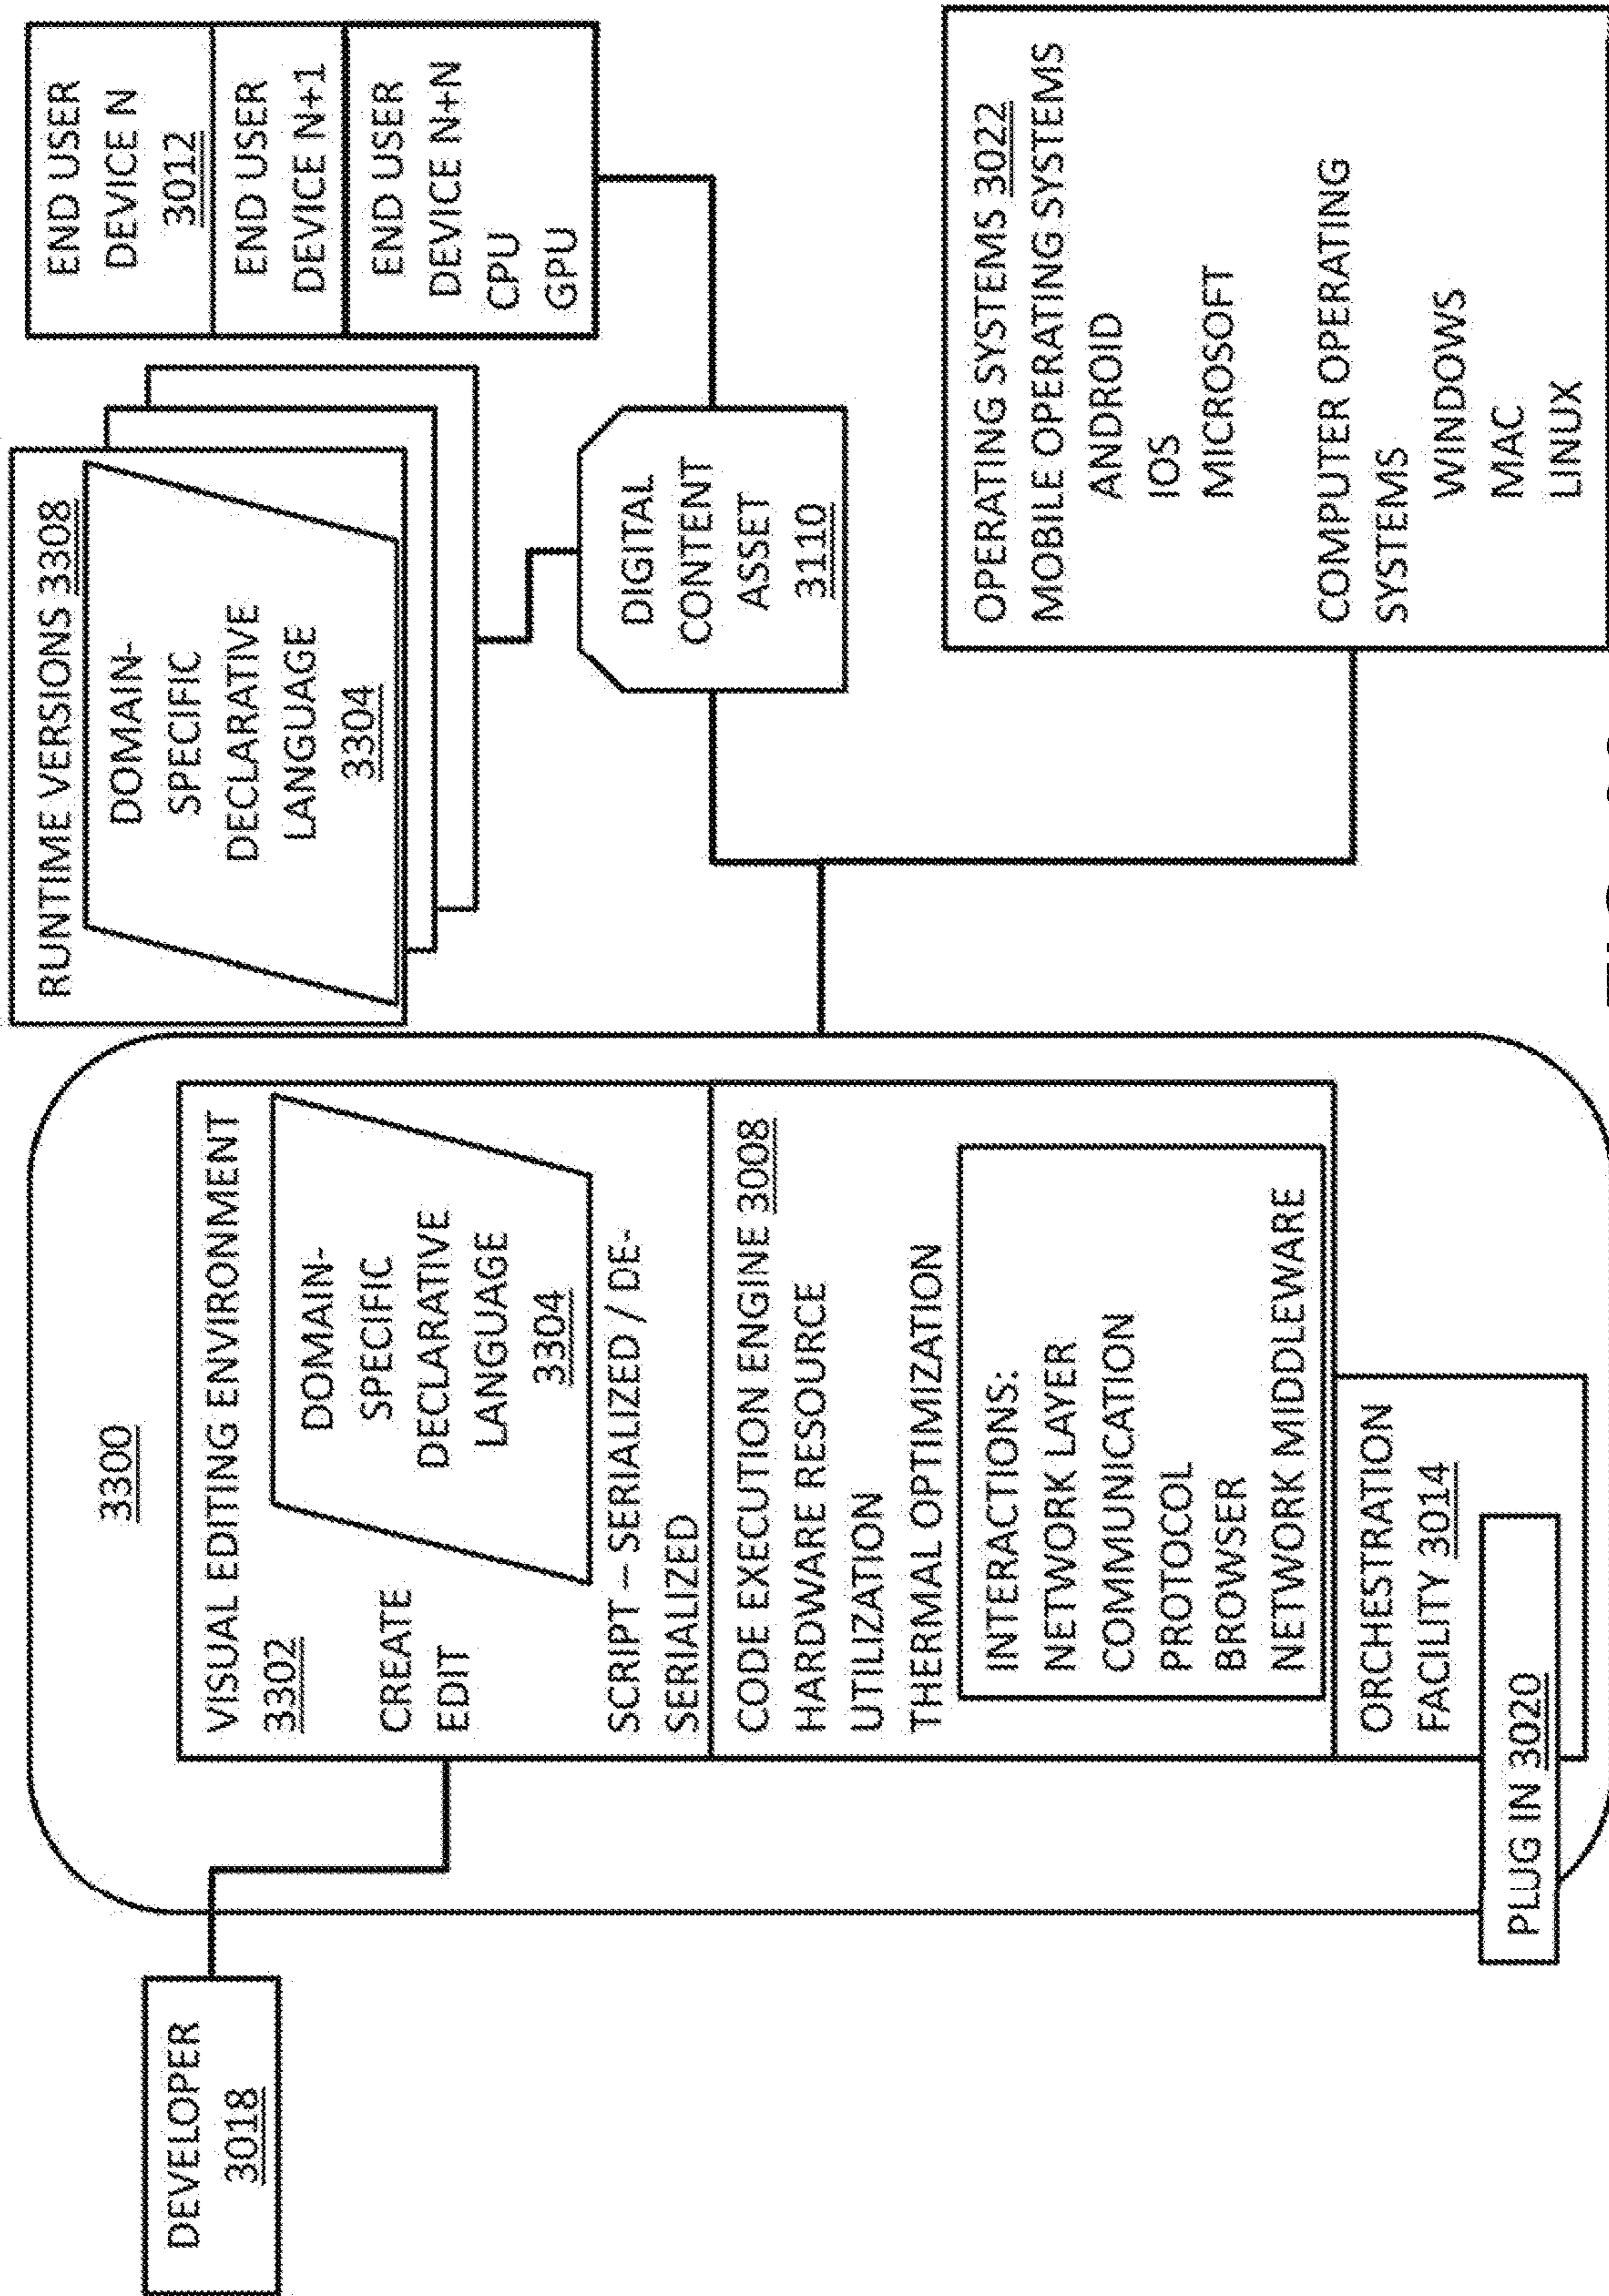


FIG. 33

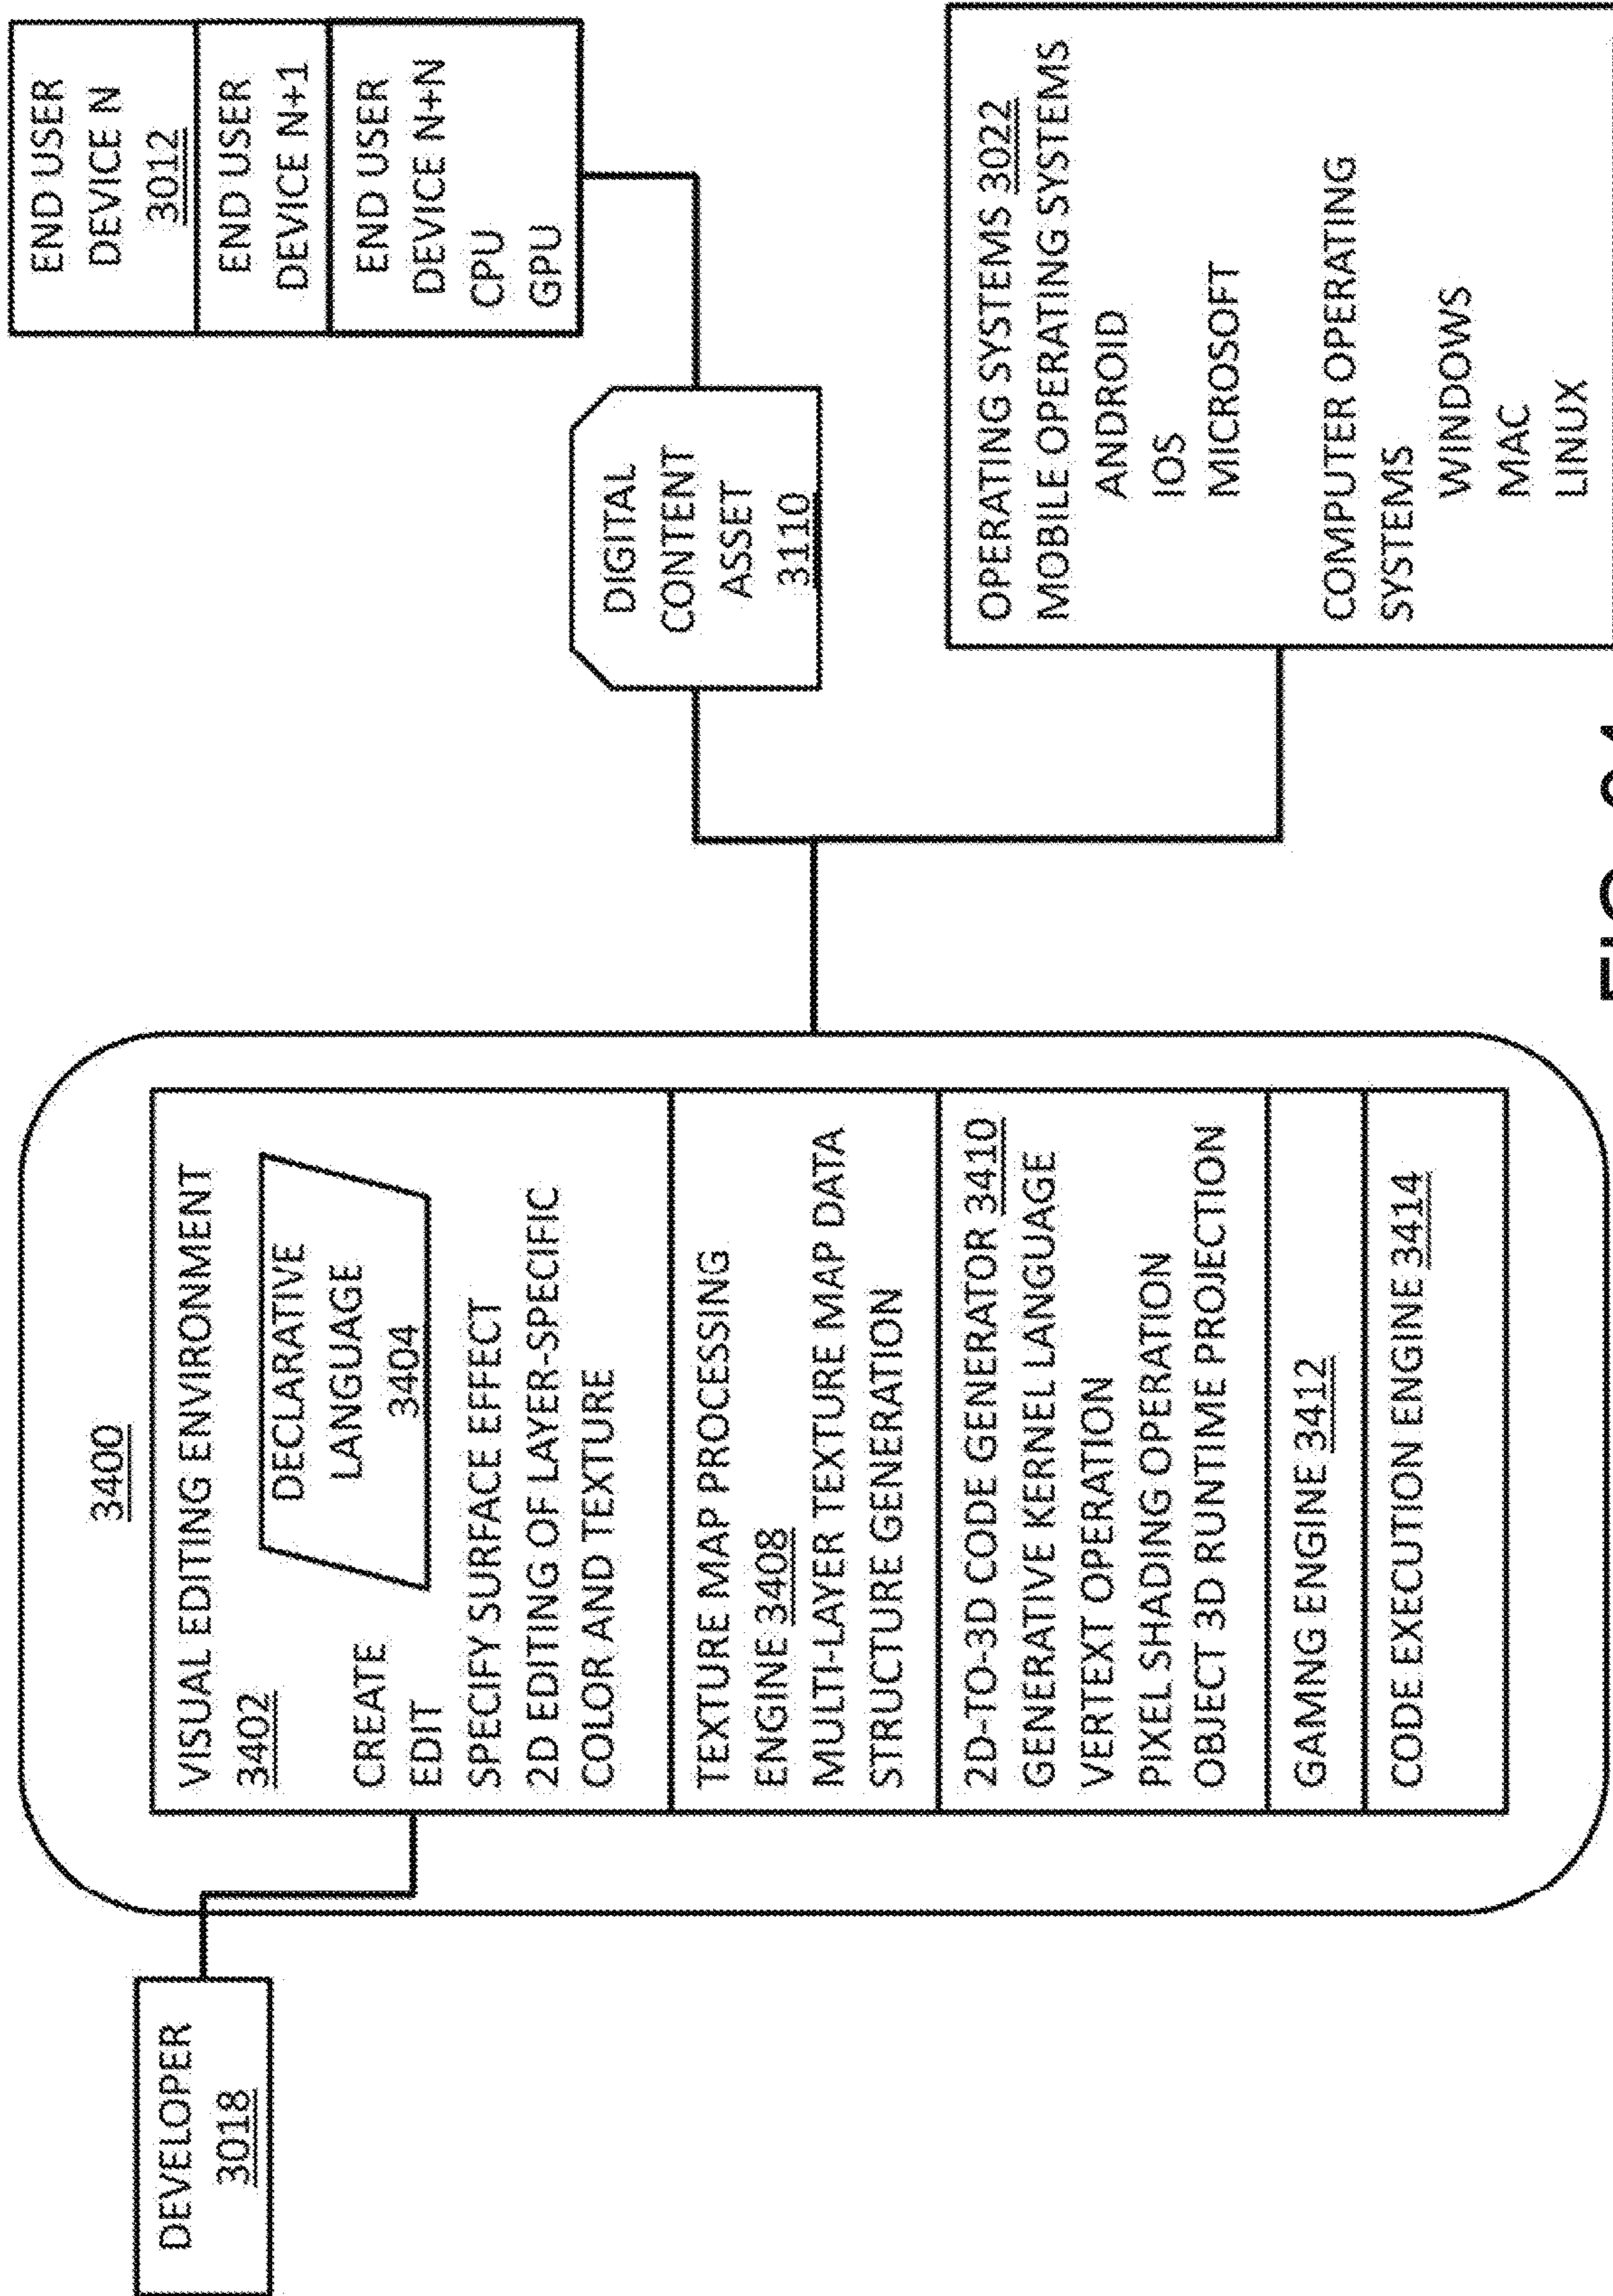


FIG. 34

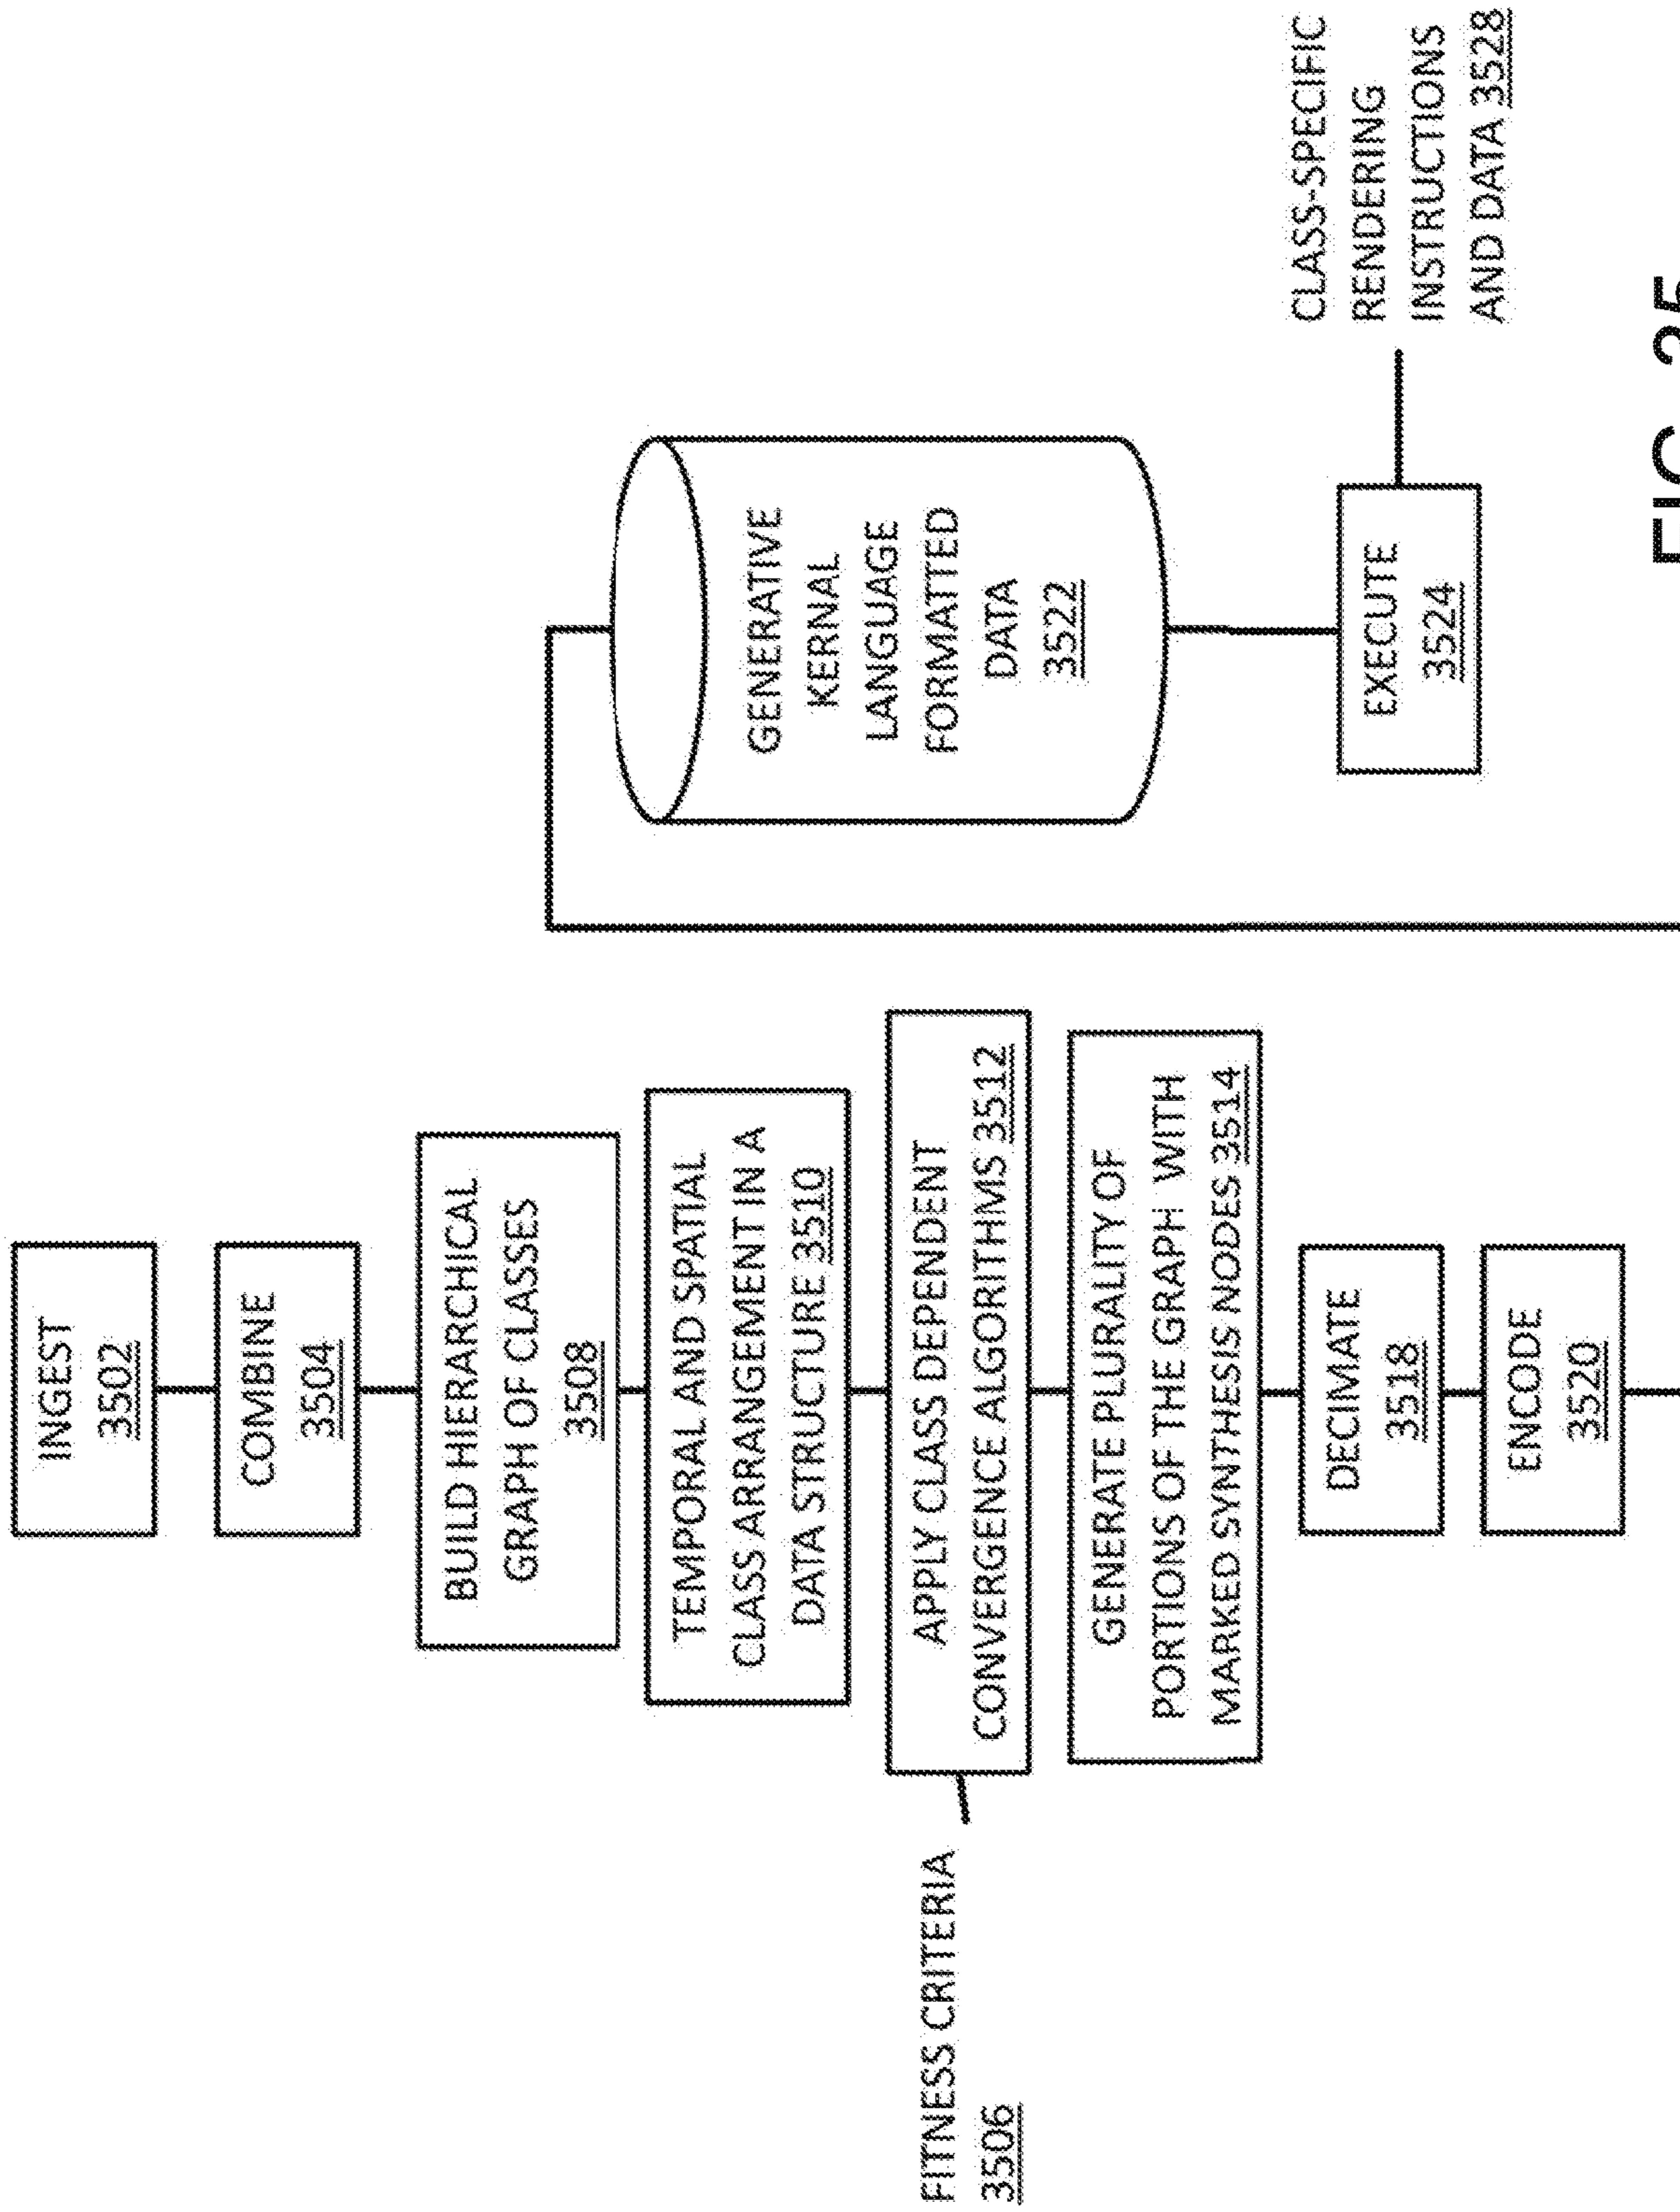


FIG. 35




FIG. 36

METHODS AND SYSTEMS FOR CREATING APPLICATIONS USING SCENE TREES

PRIORITY CLAIM

This application is a continuation of U.S. application Ser. No. 16/047,613, filed Jul. 27, 2018, entitled APPLICATION SYSTEM HAVING AN LLVM COMPILER, which is a continuation of PCT International Application No. PCT/US2018/035953, filed Jun. 5, 2018, entitled METHODS AND SYSTEMS FOR AN APPLICATION SYSTEM, which claims priority to U.S. Provisional Patent Application Ser. No. 62/515,277, filed Jun. 5, 2017, entitled METHODS AND SYSTEMS FOR A CONTENT AND DEVELOPMENT MANAGEMENT PLATFORM, U.S. Provisional Patent Application Ser. No. 62/559,940, filed Sep. 18, 2017, entitled METHODS AND SYSTEMS FOR A CONTENT AND DEVELOPMENT MANAGEMENT PLATFORM, and U.S. Provisional Patent Application Ser. No. 62/619,348, filed Jan. 19, 2018, entitled METHODS AND SYSTEMS FOR A CONTENT AND DEVELOPMENT MANAGEMENT PLATFORM, all of which are incorporated by reference as if fully set forth herein in their entirety.

BACKGROUND

1. Field

This disclosure relates to the field of software application development and digital content creation, and more particularly, to an improved technology stack for the creation, deployment, and management of computer software applications that use digital content assets.

2. Description of the Related Art

Mobile apps have become central to how enterprises and other organizations engage with both customers and employees, but many apps fall well short of customer expectations. Poor design, slow performance, inconsistent experiences across devices, and the long-time cycles and cost required in the specification, development, testing and deployment of updates requested by users are among the reasons that apps fail to engage the user or meet organizational requirements. Each type of end point device typically requires its own development effort (even with tools promising “multiplatform” support there is often the requirement for platform-specific design and coding), and relatively few app platforms can target both mobile devices and PCs while providing deep support for device capabilities like 3D, mapping, Internet of Things (IoT) integration, augmented reality (AR) and virtual reality (VR). The fact that existing app development platforms are limited in scope and require technical expertise increases time and cost while also restricting both the capability of the apps and the devices on which they can be used.

Creation, deployment and management of software applications and content that use digital content assets may be complicated, particularly when a user desires to use such applications across multiple platform types, such as involving different device types and operating systems. Software application development typically requires extensive computer coding, device and domain expertise, including knowledge of operating system behavior, knowledge of device behavior (such as how applications impact chip-level performance, battery usage, and the like) and knowledge of domain-specific languages, such that many programmers

work primarily with a given operating system type, with a given type of device, or in a given domain, and application development projects often require separate efforts, by different programmers, to port a given application from one computing environment, operating system, device type, or domain to another. While many enterprises have existing enterprise cloud systems and extensive libraries of digital content assets, such as documents, websites, logos, artwork, photographs, videos, animations, characters, music files, and many others, development projects for front end, media-rich applications that could use the assets are often highly constrained, in particular by the lack of sufficient resources that have the necessary expertise across the operating systems, devices and domains that may be involved in a given use case. As a result, most enterprises have long queues for this rich front end application development, and many applications that could serve valuable business functions are never developed because development cannot occur within the business time frame required. A need exists for methods and systems that enable rapid development of digital-content-rich applications, without requiring deep expertise in operating system behavior, device characteristics, or domain-specific programming languages. Another bottleneck is server side workflow, additional data stores and brokering of API's between existing systems.

Some efforts have been made to establish simplified programming environments that allow less sophisticated programmers to develop simple applications. The PowerApps™ service from Microsoft™ allows users within a company (with support from an IT department) to build basic mobile and web-based applications. App Maker™ by Google™ and Mobile App Builder™ from IBM™ also enable development of basic mobile and web applications. However, these platforms enable only very basic application behavior and enable development of applications for a given operating system and domain. There remains a need for a platform for developing media-rich content and applications with a simple architecture that is also comprehensive and extensible, enabling rich application behavior, low-level device control, and extensive application features, without requiring expertise in operating system behavior, expertise in device behavior, or expertise in multiple languages.

Also, it is exceedingly frustrating for a user (e.g., a developer) to continually purchase and learn multiple new development tools. A user may experience needs or requirements for debugging, error reporting, database features, image processing, handling audio or voice information, managing Internet information and search features, document viewing, localization, source code management, team management and collaboration, platform porting requirements, data transformation requirements, security requirements, and more. It may be even more difficult for a user when deploying content assets over the web or to multiple operating systems, which may generate many testing complications and difficulties. For these reasons, an application system is needed for the client that provides a full ecosystem for development and deployment of applications that work across various types of operating systems and devices without requiring platform-specific coding skills and allows for workflow, additional datastores, and API brokering on the server side.

SUMMARY

In embodiments, an application system is provided herein for enabling developers, including non-technical users, to quickly and easily create, manage, and share applications

and content that use data (such as dynamically changing enterprise data from various databases) and rich media content across personal endpoint devices. The system enables a consistent user experience for a developed application or content item across any type of device without revision or porting, including iOS and Android phones and tablets, as well as Windows, Mac, and Linux PCs.

In embodiments, the application system enables the creation, deployment and management of applications that use digital content assets. Among various methods, systems, components, services, systems, sub-systems, and other elements, the application system may include an engine that supports development of an application by using a visual editor, where the same engine supports runtime operation of the application. In embodiments, the system uses a novel declarative language that includes layout and logic capabilities (and is used both as a file format and for producing binary machine code to control low-level device behavior). The language enables rich application behavior using simple, intuitive commands that can be readily mastered by non-experts. The system enables publication through a viewer/portal, where application and content behavior are contained in a container that enables convenient publication and consumption across endpoint devices and platforms. The engine includes a wide range of capabilities, features, and extensions for enabling a wide range of content and application behaviors and capabilities, such as for enabling avatars, 3D geometry features, cinematic and animated presentation effects, and many others as are described throughout this disclosure. A server element which allows for workflow, additional data storage, API brokering, and for the hosting of published applications is also disclosed.

In embodiments, the application system includes an engine that has many features typically deployed in engines for high-performance video games, such that the system enables extremely fast application behavior and has built-in support for advanced graphics (including 3D geometry and physics for animation), as well as capabilities for handling machine vision, augmented and virtual reality, online and offline maps, and gesture control. An optionally multi-user visual editor of the system allows dynamic data from any source (such as documents, enterprise databases, media files, internet feeds, IoT cloud databases etc.) to be published in a container that will run in a viewer across any end point device and/or can be published to a public app store. Content updates and variations in published content or applications, such as by customer segment or geography, can be easily added by content owners.

In embodiments, the application system described herein fundamentally changes the process of delivering information to endpoint devices and creating an engaging interactive experience. The visual editor allows multiple groups to simultaneously lay out and structure information that is packaged into a container for publication. This can be done in minutes for simple projects, and in no more than a few days for projects that draw extensively on external information feeds and utilize complex graphics. The container with developed application and content in it can be instantly published, such as to a cloud platform (such as the Amazon Web Services™ cloud), where it can then be immediately accessed by authorized users, on any device, with an identical experience regardless of the device or platform, or to a public or private app store, where it can be available as soon as any other requirements for publication in the particular app store are satisfied.

In embodiments, the application system, language and overall implementation are simple in that there is only a

limited degree of complexity and there are few layers of abstraction. Simple, in this description, does not indicate the opposite of powerful. For example, the system, declarative language, and overall implementation may minimize syntax, with no case sensitivity; provide a highly consistent way of doing things; support well understood media formats; use a small runtime; support code sharing and collaboration; and use the same language for the editor, viewer and portal. A complete application system may seem mutually exclusive to being simple, but it is not. The application system disclosed herein covers the domain of tools required to build and deploy solutions, while maintaining integration and simplicity.

Because of these complete functions, in embodiments an application system may be open at design time and closed at runtime. This is an important distinction, because other developers may be free to build tools and export applications from the application system from their products for design time, but at runtime, the environment may be closed, with the function provided by the common engine editor and runtime infrastructure of the application system providing security and stability. A custom runtime with additional platform-specific binary code linked in can be built, but the security and stability cannot necessarily be guaranteed.

In embodiments, the engine may allow a user to quickly and easily create, manage, and share enterprise information and rich media content across personal endpoint devices. The engine may include an editor, a runtime interface and a user interface, all built on top of the engine.

In embodiments, the visual editor may include the visual editor, which may be part of a visual editing environment with a common engine for editing and runtime, which in embodiments may provide a real-time, multi-user simultaneous development environment, such as for multiple developers. The visual editor may allow an engine full benefit of managing utilization of a CPU, a GPU, or the like, in a visual editing environment for software code development. In embodiments, the visual editing environment may include capabilities for a gaming engine, including for coding gaming behavior of hardware and software components, including operating system and hardware platforms.

The system may also involve a cluster of technologies and capabilities related to a unique, dynamic, declarative language. This “dynamic language cluster” may include capabilities for creating and handling a scene tree. In embodiments, the editor and runtime for the dynamic language may be compiled with a compiler (such as the LLVM™ compiler). The language may enable the ability to express logic and layout in the same language. The language may use the same domain-specific language for creating the editor/runtime and for use as a file format. The language may include state information built into the language and expose modules. The language may support creation of visual experiences and allow a user to publish apps through a private portal. In embodiments the language may be configured to be processed linearly (i.e., without loops), such as by processing a directed acyclic graph or similar structure.

An avatar engine may support avatar functions, such as used in commercial user interfaces, conversational interfaces and natural language endpoints.

In embodiments, a cluster of technologies and components relating to lower level code, referred to as the “low code technology cluster,” may support variations extended with JavaScript. Variations may be accessed via a variations user interface.

Among other elements of the application system, an interface and supporting processing elements and systems

5

may be provided for developing applications that render three-dimensional visual effects and that handle three-dimensional objects and their geometric parameters, according to some embodiments. These elements are referred to herein for simplicity in some cases as the three-dimensional user interface, or simply the 3D UI. The interface may support input, display and manipulation of 3D objects. In embodiments, the 3D UI may render apps in stereo with 3D input for virtual reality (VR) and augmented reality (AR). The 3D UI may include interactions from a range of inputs including 3D machine vision, which in embodiments may include AR experiences. The 3D UI may mix 2D and 3D elements in the same application, such as, in embodiments, using a hybrid scene tree system.

In embodiments, a system is provided, including a visual editor for creating and editing an application in declarative language, wherein the language provides the capability for specifying control of machine behavior of a device, the abstraction of input types across operating system types and the capability for control of visual presentation layer behavior of objects across a plurality of operating system platform types; and an engine having gaming engine features, wherein the editor and the runtime of an application are created using the system.

In embodiments, a system for creating, sharing and managing digital content may include a visual editing environment that enables a developer to create and edit code controlling a digital content asset in a declarative language and a code execution engine that operates on the code created in the visual editing environment to control execution of at least one hardware infrastructure element that enables the utilization of the digital content asset, wherein the same declarative language and code execution engine are used during editing of the digital content asset and at runtime and wherein the visual editor and runtime code are compiled for distribution using LLVM. In embodiments, the visual editor and runtime code are bound to the code execution engine. In embodiments, the code execution engine is a C++ engine. In embodiments, a plurality of end users using different devices can simultaneously experience the same behavior of the digital content asset during its creation and editing. In embodiments, a digital content asset may be compiled, linked to the code execution engine and executed by the engine independent of support by tools outside the visual editor and code execution engine. In embodiments, the code execution engine controls at least one of CPU and GPU utilization for a hardware endpoint device upon which the digital content asset runs. In embodiments, the code execution engine governs execution of the code across a plurality of operating systems to provide a consistent user experience with the digital content asset. In embodiments, the code execution engine governs execution of the code across a plurality of mobile operating systems. In embodiments, the mobile operating systems are selected from among IOS, Android and Windows operating systems. In embodiments, the code execution engine governs execution of the code across a plurality of computer operating systems. In embodiments, the computer systems are selected from among Windows, Mac, and Linux operating systems. In embodiments, the code execution engine governs execution of the code across a plurality of operating systems including at least one mobile operating system and at least one other operating system. In embodiments, the operating systems are selected from the group consisting of Windows, Mac, Linux, IOS and Android operating systems. In embodiments, the code execution engine enables control of network layer interactions for the digital content asset. In embodi-

6

ments, the system may include at least one LLVM debugging tool for debugging the compiled code for the digital content asset. In embodiments, the system may include a plug-in system for orchestrating components and events for the digital content asset. In embodiments, the plug-in system is a JavaScript plug-in system. In embodiments, the visual editing environment is configured to serialize code input to the visual editing environment. In embodiments, the visual editing environment is configured to de-serialize code input to the visual editing environment.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1A provides a schematic diagram of the main components of an application system and interactions among the components in accordance with the many embodiments of the present disclosure.

FIG. 1B depicts a schematic diagram of a server kit ecosystem of an application system in accordance with the many embodiments of the present disclosure.

FIG. 1C depicts a schematic diagram of a server kit software appliance of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 2A, 2B, and 2C depict screenshots of embodiments of an editor of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 2D, 2E, and 2F depict a collaborative 3D workspace of an application system in accordance with the many embodiments of the present disclosure.

FIG. 2G depicts a three-stage example of a high-level recipe workflow of an application system in accordance with the many embodiments of the present disclosure.

FIG. 3 depicts a continuum of natural user interface interaction which a user may use to interact with content created by an application system in accordance with the many embodiments of the present disclosure.

FIG. 4 depicts a project built by an app builder of an application system in accordance with the many embodiments of the present disclosure.

FIG. 5 depicts an instance example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 6 depicts a nested instance example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 7A and FIG. 7B depict a define-and-raise function example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 8 depicts an expression example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 9A depicts a simple “if” block condition example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 9B depicts an “if/else if/else” block condition example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 10A depicts a simple “loop if” pre-tested loop example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 10B depicts a simple “end if” posted test loop example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 10C depicts an iterator-style “for loop” example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 10D depicts an array/list or map iterator-style “collection” loop example of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 11A-11D depict array examples of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 12A and 12B depict variation and state examples of an application system in accordance with the many embodiments of the present disclosure.

FIG. 13 depicts example variation rules of an application system in accordance with the many embodiments of the present disclosure.

FIG. 14 depicts the declarative language scene tree description of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 15A and 15B depict an example of a button specified using conditional logic of an application system in accordance with the many embodiments of the present disclosure.

FIG. 16 depicts a scene tree description example of an SGML element nesting example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 17 depicts an example of logic placement inside a method of an application system in accordance with the many embodiments of the present disclosure.

FIG. 18 depicts an example of how statically declared states may be used to determine unpicked and picked visualizations of an application system in accordance with the many embodiments of the present disclosure.

FIG. 19 depicts an avatar character example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 20 depicts an avatar morph target example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 21 depicts an avatar car interaction example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 22 depicts an avatar gaze and direction example of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 23A and 23B depict morph vertex animation examples of an application system in accordance with the many embodiments of the present disclosure.

FIG. 24 depicts a content density example of an application system in accordance with the many embodiments of the present disclosure.

FIG. 25 depicts a partial hand view of a user interacting with content in accordance with the many embodiments of the present disclosure.

FIGS. 26A and 26B depict hand centered approaches of an application system 100 in accordance with the many embodiments of the present disclosure.

FIGS. 27A and 27B depict examples of a user interacting with content by closing their hands into a fist in accordance with the many embodiments of the present disclosure.

FIGS. 28A-28D depict halo fingertip detection examples of an application system in accordance with the many embodiments of the present disclosure.

FIGS. 29A and 29B depict halo fingertip detection examples of an application system in accordance with the many embodiments of the present disclosure.

FIG. 30 depicts a system diagram of an embodiment of a system for creating, sharing and managing digital content in accordance with the many embodiments of the present disclosure.

FIG. 31 depicts a system diagram of an embodiment of a system for creating, sharing and managing digital content with a runtime that shares a declarative language with a visual editing environment in accordance with the many embodiments of the present disclosure.

FIG. 32 depicts a system diagram of an embodiment of a system for creating, sharing and managing digital content with a gaming engine capability in accordance with the many embodiments of the present disclosure.

FIG. 33 depicts a system diagram of an embodiment of a system for creating, sharing and managing digital content with a plurality of runtimes that share a domain-specific declarative language with a visual editing environment in accordance with the many embodiments of the present disclosure.

FIG. 34 depicts a system diagram of an embodiment of a system for creating, sharing and managing digital content that includes texture mapping and 2D-to-3D code generation in accordance with the many embodiments of the present disclosure.

FIG. 35 depicts a flow chart of an embodiment of producing generative content with use of a generative kernel language in accordance with the many embodiments of the present disclosure.

FIG. 36 depicts a system diagram of an embodiment of an application system environment engine in accordance with the many embodiments of the present disclosure.

DETAILED DESCRIPTION

FIG. 1A depicts embodiments of an application system 100 (also referred to as a “content and development management platform” or “platform 100” in the incorporated materials) according to exemplary and non-limiting embodiments. In embodiments, the application system 100 may include an engine 102 and an editor and runtime infrastructure 104 that includes various associated components, services, systems, and the like for creating and publishing content and applications 178. These may include a content and application creator 109 (referred to in some cases herein for simplicity of reference as the app creator 109). In embodiments, the app creator 109 may include a visual editor 108 and other components for creating content and applications 178 and a declarative language 140 which has hierarchical properties, layout properties, static language properties (such as properties found in strongly typed and fully compiled languages) and dynamic language properties (e.g., in embodiments including runtime introspection, scripting properties and other properties typically found in non-compiled languages like JavaScript™). In embodiments, the declarative language 140 is a language with simple syntax but very powerful capabilities as described elsewhere in this disclosure and is provided as part of a stack of components for coding, compiling and publishing applications and content. In embodiments the declarative language has a syntax that is configured to be parsed linearly (i.e., without loops), such as taking the form of a directed acyclic graph. Also included is a publisher 138, which in embodiments includes a compiler 136. In embodiments, the compiler is an LLVM compiler 136. The publisher 138 may further include other components for compiling and publishing content and applications. Various components of the editor and runtime infrastructure 104 may use the same engine 102, such as by integrating directly with engine components or by accessing one or more application programming interfaces, collectively referred to herein as the engine API 114.

The application system **100** may be designed to support the development of media-rich content and applications. The application system **100** is intended to have a simple architecture, while providing comprehensive and extensible functionality.

The language used by the application system **100** may, in embodiments, have a limited degree of flexibility and few layers of abstraction, while still providing a broad and powerful set of functions to a user. In embodiments, characteristics of the application system **100** with limited flexibility may include minimizing syntax requirements, such as by excluding case sensitivity, supporting a single way of completing tasks, without exceptions, supporting a limited number of media formats, helping keep the runtime small, sharing code and using a common language for editor, viewer and portal.

While an application system **100** may have a limited degree of flexibility and few layers of abstraction, it may provide a comprehensive range of capabilities required to build and deploy solutions. This comprehensive range of capabilities may be provided and performed while maintaining an integrated and simple, i.e., straightforward, mode of operation for a user.

In embodiments, the engine **102** may include a wide range of features, components and capabilities that are typical of engines normally used for high performance video games, including, but not limited to, an avatar interaction and rendering engine **148** (referred to for simplicity in some cases as simply the avatar engine **148**), a gaming engine **119**, a physics engine **152**, a virtual reality (VR) engine **154**, an augmented reality (AR) engine **158**, a machine vision engine **160**, an animation engine **162**, a 3D engine **164**, a rich user interaction engine **168** (such as for handling gesture, touch, and voice interaction), a geographic information system (GIS) way-finding and map engine **170**, audio, video and image effects engine **194**, and/or e-commerce engine **182**. These may be distinct components or systems that interact with or connect to each other and to other system components, or they may be integrated with each other. For example, in embodiments, the animation engine **162** may include capabilities for handling 3D visual effects, physics and geometry of objects, or those elements may be handled by a separate system or engine.

The engine **102** may also include a browser engine **226**. In embodiments, the browser engine **226** may be a lightweight JavaScript implementation of a subset of the engine **102**. The browser engine **226** may render on a per-frame basis inside a web browser, using WebGL 2.0, for example, to render without restrictions, such as restrictions that may be imposed by rendering in HTML 5.0.

In embodiments, rendering may be a two-step process. A first step may use asm.js, a subset of JavaScript that is designed to execute quickly, without using dynamic features of JavaScript. Asm.js may be targeted by an LLVM compiler **136** of the application system **100**. This may allow an application of the application system **100** to be compiled in asm.js.

In a second step, a C++ engine and OS level wrappers may be created to target web browser capabilities and also use WebGL 2.0 for rendering. This subset may significantly reduce the libraries required to render, as they may be replaced with equivalent capabilities, such as those provided by a modern browser with network and sound support. The target may be a small set of C++ engine code, which may also be compiled by the application system **100**, for example

by the LLVM compiler **136**, to target asm.js and become the engine that applications of the application system **100** may run against.

The two-step process may produce high performance, sovereign applications rendered in WebGL that can operate inside a browser window. Because the asm.js files of application system **100** may be the same for all applications of this type, they may be cached, allowing improved start-up times across multiple applications of the application system **100**. This approach may remove the limitations caused by HTML and CSS, as well as memory usage and rendering performance issues that are experienced by current state-of-the-art websites.

The application system **100** may include various internal and external communication facilities **184**, such as using various networking and software communication protocols (including network interfaces, application programming interfaces, database interfaces, search capabilities, and the like). The application system **100** may include tools for encryption, security, access control and the like. The application system **100** may consume and use data from various data sources **172**, such as from enterprise databases, such that applications and content provided by the system may dynamically change in response to changes in the data sources **172**. Data sources **172** may include documents, enterprise databases, media files, Internet feeds, data from devices (such as within the Internet of Things), cloud databases, and many others. The application system **100** may connect to and interact with and various cloud services **142**, third party modules **146**, and applications (including content and applications **178** developed with the application system **100** and other applications **150** that may provide content to or consume content from the application system **100**). Within the editor **108** and language **140**, and enabled by the engine **102**, the application system **100** may allow for control of low level device behavior, such for how an endpoint device will render or execute an application, such as providing control of processor utilization **188** (e.g., CPU and or GPU utilization). The application system **100** may have one or more plug-in systems, such as a JavaScript plug-in system, for using plug-ins or taking input from external systems.

The application system **100** may be used to develop content and applications **178**, such as ones that have animated, cinematic visual effects and that reflect, such as in application behavior, dynamic changes in the data sources **172**. The content and applications **178** may be deployed in the viewer/portal **144** that enables viewing and interaction, with a consistent user experience, upon various endpoint devices (such as mobile phones, tablets, personal computers, and the like). The application system **100** may enable creation of variations **190** of a given content item or application **178**, such as for localization of content to a particular geographic region, customization to a particular domain, user group, or the like. Content and applications **178** (referred to for simplicity in some cases herein as simply applications **178**), including any applicable variations **190**, may be deployed in a container **174** that can run in the viewer and that allows them to be published, such as to the cloud (such as through a cloud platform like Amazon Web Services™) or via any private or public app store.

A user interface **106** may provide access, such as by a developer, to the functionality of the various components of the application system **100**, including via the visual editor **108**. In embodiments, the user interface **106** may be a unified interface providing access to all elements of the system **100**, or it may comprise a set of user interfaces **106**, each of which

11

provides access to one or more elements of the engine **102** and other elements of the application system **100**.

In embodiments, the application system **100** may support the novel declarative language **140** discussed herein. The application system **100** may support additional or alternative programming languages as well, including other declarative languages. A cluster of technologies around the declarative language **140** may include a publisher **138** for publishing content and applications (such as in the container **174**) and a compiler (e.g., an LLVM compiler **136**). The LLVM compiler **136** may comprise one or more libraries that may be used to construct, optimize and produce intermediate and/or binary machine code from the front-end code developed using the declarative language **140**. The declarative language **140** may include various features, classes, objects, functions and parameters that enable simple, powerful creation of applications that render assets with powerful visual effects. These may include domain-specific scripts **134**, a scene tree description system **124**, logic **126**, a file format system **128** for handling various file formats, and/or an object state information system **130** (also referred to as “state” in this disclosure). The application system **100** may include various other components, features, services, plugins and modules (collectively referred to herein as modules **132**), such as for engaging the various capabilities of the engine **102** or other capabilities of the application system **100**. The declarative language **140** and surrounding cluster of technologies may connect to and operate with various application system **100**, such as the engine **102**, third party modules **146**, cloud services **142** and the visual editor **108**.

The visual editor **108** of the application system **100** may be designed to facilitate rapid creation of content, optionally allowing users to draw from an extensive set of templates and blueprints **180** (which may be stored in the data store **172**) that allow non-technical users to create simple apps and content, much as they would create a slide show or presentation, such as using a desktop application like Microsoft™ PowerPoint™.

In embodiments, the visual editor **108** may interact with the engine **102**, (e.g., via an engine application programming interface (API) **114**), and may include, connect to, or integrate with an abstraction engine **118**, a runtime editor **110**, a serialization engine **112**, and/or a capability for collaboration and synchronization of applications and assets in a multi-user development environment **120** (referred to for simplicity in some cases as “multi-user sync”). The multi-user sync system **120** may operate as part of the editor and runtime infrastructure **104** to allow simultaneous editing by multiple users, such as through multiple instances of the visual editor **108**. In this way, multiple users may contribute to an application coextensively and/or, as discussed below, may view a simulation of the application in real time, without a need for compilation and deployment of the application.

In embodiments, the visual editor **108** may allow editing of code written in the declarative language **140** while displaying visual content that reflects the current state of the application or content being edited **116** in the user interface **106** (such as visual effects). In this way, a user can undertake coding and immediately see the effects of such coding on the same screen. This may include simultaneous viewing by multiple users, who may see edits made by other users and see the effects created by the edits in real time. The visual editor **108** may connect to and enable elements from cloud services **142**, third party modules **146**, and the engine **102**. The visual editor **108** may include an abstraction engine **118**,

12

such as for handling abstraction of lower-level code to higher-level objects and functions.

In embodiments, the visual editor **108** may provide a full 3D rendering engine for text, images, animation, maps, models and other similar content types. In embodiments, a full 3D rendering engine may allow a user (e.g., developer) to create, preview, and test 3D content (e.g., a virtual reality application or 3D video game) in real time. The visual editor **108** may also decompose traditional code driven elements, which may comprise simplified procedural logic presented as a simple list of actions, instead of converting them into a user interface, which may comprise simplified conditional logic presented as a visual checklist of exceptions called variations.

In embodiments, editing code in the declarative language **140** within the visual editor **108** may enable capabilities of the various engines, components, features, capabilities and systems of the engine **102**, such as gaming engine features, avatar features, gestural interface features, realistic, animated behavior of objects (such as following rules of physics and geometry within 2D and 3D environments), AR and VR features, map-based features, and many others.

The elements of the application system **100** described in the foregoing paragraphs and the other elements described throughout this disclosure may connect to or be integrated with each other in various configurations to enable the capabilities described herein.

As noted above, the application system **100** may include a number of elements, including the engine **102**, a viewer **144** (and/or portal), an application creator (e.g., using the visual editor **108** and declarative language **140**), and various other elements, including cloud services **142**.

In embodiments, the engine **102** may be a C++ engine, which may be compiled and may provide an operating system (OS) layer and core hardware accelerated functionality. The engine **102** may be bound with LLVM to provide just-in-time (JIT) compiling of a domain-specific script **134**. In embodiments, the LLVM compiler **136** may be configured to fully pre-compile the application to intermediate ‘bytecodes’ or to binary code on-demand. In embodiments, the LLVM compiler **136** may be configured to activate when a method is called and may compile bytecodes of just this method into native machine code, where the compiling occurs “just-in-time” to run on the applicable machine. When a method has been compiled, a machine (including a virtual machine) can call the compiled code of the method, rather than requiring it to be interpreted. The engine **102** may also be used as part of a tool chain along with the LLVM compiler **136**, which may avoid the need to provide extra code with a final application.

In embodiments, the design of the underlying C++ engine of the application system **100** may be built around a multi-threaded game style engine. This multi-threaded game style engine may marshal resources, cache media and data, manage textures, and handle sound, network traffic, animation features, physics for moving objects and shaders.

At the center of these processes may be a high performance shared scene tree. Such a scene tree is non-trivial, such as for multi-threading, and it may be serialized into the language of the application system **100**. Doing this may allow the scene tree to act as an object with properties and actions associated with events. There may then be modular layers for managing shared implementation information in C++, as well as platform-specific implementations for these objects in the appropriate language, such as objective C or Java and also allowing binding to the appropriate API’s or SDK’s or Libraries.

As well as serializing the project in the application system format, it may also be possible to export a scene tree as JSON or in a binary format. Additionally, only a subset of the full application system language may be required for the editor **108** and viewer/portal **144**. Such a subset may support objects/components, properties, states and lists of method calls against events. The subset may also be suitable for exporting to JSON. The scene tree may also be provided as a low level binary format, which may explicitly define the structures, data types, and/or lengths of variable length data, of all records and values written out. This may provide extremely fast loading and saving, as there is no parsing phase at all. The ability to serialize to other formats may also make it more efficient for porting data to other operating systems and software containers, such as the Java Virtual Machine and Runtime (JVM) or an asm.js framework inside a WebGL capable browser.

In many cases the application system **100** may not have direct access to device hardware for devices running other operating systems or software containers, so various things may be unknown to the application system **100**. For example, the application system **100** may not know what libraries are present on the devices, what filtering is used for rendering quads, what font rasterizer is used, and the like. However, the system's ability to produce or work with "compatible" viewers, which may exist inside another technology stack (e.g., in the Java runtime or a modern web browser), may allow users to experience the same project produced using the application system **100** on an alternative software platform of choice. This may also provide a developer of the viewer the same opportunity to build in similar game engine-style optimizations that may take effect under the hood, as may be possible within the application system **100** itself.

The viewer **144** (also referred to as a "portal" or "viewer/editor/portal" in this disclosure) may integrate with or be a companion application to a main content and application creator **109**, such as one having the visual editor **108**. The viewer **144** may be able to view applications created without the need to edit them. The ability to load these applications without the need for binary compilation (e.g., by LLVM) may allow applications to run with data supplied to them. For example, objects may be serialized into memory, and built-in functions, along with any included JavaScript functions, may be triggered.

The app creator **109**, also referred to in some cases as the visual editor **108** in this disclosure, may be built using the engine **102** of the application system **100** and may enable editing in the declarative language **140**. FIGS. 2A-2C depict screenshots of an editor. An app creator may allow a user to take advantage of the power of an application system **100** engine **102**.

The various editors of the application system **100** may allow a user to effectively edit live application system **100** objects inside of a sandbox (e.g., a contained environment). In embodiments, the editors (e.g., runtime editor **110** and/or visual editor **108**) may take advantage of the application system **100** ability to be reflective and serialize objects to and from the sandbox. This may have huge benefits for simplicity and may allow users to experience the same behavior within the editor as they do when a project (e.g., an application) deploys because the same engine **102** hosts the editors, along with the project being developed **108** and the publisher **138** of the editor and runtime infrastructure **104**.

In embodiments, the application system **100** may exploit several important concepts to make the process of develop-

ment much easier and to move the partition where writing code would normally be required.

The application system **100** may provide a full layout with a powerful embedded animation system, which may have unique properties. The application system **100** may decompose traditional code driven elements using simple linear action sequences (simplified procedural logic), with the variations system **190** to create variations of an application **178**, optionally using a visual checklist of exceptions (simplified conditional logic), and the like.

In embodiments, the application system **100** may break the creation of large applications into projects. Projects may include pages, components, actions and feeds. Pages may hold components. Components may be visual building blocks of a layout. Actions may be procedural logic commands which may be associated with component events. Feeds may be feeds of data associated with component properties. Projects may be extended by developers with JavaScript plugins. This may allow third party modules or plugins **146** to have access to all the power of the engine. The engine **102** may perform most of the compute-intensive processing, while JavaScript may be used to configure the behavior of the engine.

In embodiments, the cloud services **142** of the application system **100** may be a series of important server side elements including the compiler wall (which builds final binaries), file versioning (such as using an API similar to what is used with simplified, basic cloud storage systems, such as the Simple Storage Service (S3™) from Amazon™), analytics capture and display, content editing and distribution, project distribution, a real-time communications hub and other related server-side elements. As depicted in FIG. 1B, a cloud services architecture **142** may include a server kit software appliance or server kit **186**. The server kit **186** may be a software appliance that may be installed on the cloud services architecture **142** of a client, such as a Microsoft Azure™ or Amazon AWS™ cloud services architecture **142**, to aid in creating a turnkey secure and scalable backend REST and Real Time APIs. The server kit **186** may allow mobile apps or web apps to connect to required enterprise IT infrastructure elements, represented by content and applications **178** in FIG. 1B, in a secure and system agnostic manner. The server kit **186** may be configured without coding, such as using a web GUI that may edit a configuration file.

The server kit **186** may be provisioned or installed from a marketplace/store associated with the selected cloud services architecture and may connect to applications **150** of the application system **100**. In embodiments, the server kit **186** may provide a web console that may allow the configuration of various different connections and settings. In embodiments, all aspects of configuration may be data driven. A configuration may include multiple aggregate options to acquire user data.

For users and groups of users, it may possible to set the access rights to heterogeneous assets, such as database views, REST calls, files and folders, that may then be inside the client firewalled environment, such as for providing security. These assets may be set up as live connections, or they may be polled with a nominated frequency. The server kit **186** may also include the ability to redirect to cache database read views, REST responses and files, such as to reduce load on internal systems and improve download performance through the Internet's existing proxy architecture. When assets are added or updated, all applications connected may be updated using live broadcast.

15

In embodiments, a server kit **186** may provide a simple workflow system that may enable chaining incoming requests and one or multiple server side processes, to one or multiple heterogeneous systems. A workflow system may keep all applications connected updated of changes in state using a live broadcast. Workflow nodes may use basic configurable logic or custom JS.

In embodiments, a server kit **186** may provide a high performance, secure backend API for web and native applications. A server kit **186** may provide this by using caching where possible, broadcasts and responses to real time events from many different underlying enterprise systems and resources that may not need to be aware of each other. In embodiments, a server kit **186** may provide a securely managed, flexible & loosely coupled layer on top of an enterprise ecosystem. The server kit **186** may expose only the elements required and only to authorized apps and authorized users of those apps.

As depicted in FIG. 1C, an example server kit **186** may form a combination of core capabilities, as well as some underlying features within its modules that make the overall capability possible. The configuration of the server kit **186** may be data driven. It may only be when a user determines a need for more capability that the user may use either JavaScript or may specify a REST hook to override the provided systems for a workflow rule or a data transformation. In embodiments, the server kit **186** may include a data transformation capability **188**. The data transformation capability **188** may include REST data ingestion capability **190**. The REST data ingestion capability **190** may use XSLT for XML and JSONT for incoming JSON data, to transform the data for storage and future use or custom transformation specified as JS nodes.

In embodiments, a server kit **186** may include a data export capability. The data export capability may include, for example, structured templating, free-flow text templating (in HTML for example), and visual templating (e.g., for PDF, PNG and JPG), and may be provided to render out assets from data. Fully custom rules may be provided with JS nodes.

In embodiments, a server kit **186** may include support for clonability. Importantly, the server kit **186** and its related data-structures may be configured to be able to be cloned. The server kit **186** and its related data-structures may be cloned without requiring data duplication on the enterprise side. However, items inside an appliance and nominated edge cache may end up the same when requested.

In embodiments, a server kit **186** may include support for cascaded inserts. Cascaded inserts may support a cascading view **192**, which may be a single compressed view submission, to populate both a header record and parameterize sub-queries in a cascade, to populate child record.

A task-based workflow for an application system **100** may hold lists of items as it proceeds through the execution of each task. Items in these lists may include the state of each item, the user who owns each item, the basic rules that determine changing state and the actions performed when that rule is triggered. All of these items may be configurable in a Web GUI of an application system **100** without coding, while also including support for custom rules to make external REST calls or run supplied JS functions, which may be processed by JS rule **194**. In embodiments, the server kit **186** may trigger workflows based on server alerts. Workflows may be triggered and managed using a workflow system **198**.

In embodiments, a server kit **186** may include configuration plugins. The administrator of the server kit **186** may

16

activate a plugin that may provide a default configuration along with a schema file provided by a developer. The default configuration may then configure the nominated schema on, for example, the SQL system chosen by an administrator, including data that was initially specified.

In embodiments, a server kit **186** may issue and revoke tokens for applications **150** and see analytics **196** on application usage and transaction logs. Applications **150** may be web and native applications.

In embodiments, a server kit **186** may include users **198** and assign rights **200** and roles **202** to users **198**. Rights **200** and roles **202** may include access to REST calls, DSTORE functions and the like. Tasks may be assigned a user and state.

In embodiments, a server kit **186** may also support the ability to transform data using JSONT, cascading inserts, pivot table queries and mail merge. In exemplary and non-limiting embodiments, the server kit **186** may support mail merge using a PHP library to generate HTML code.

In embodiments, a server kit **186** may also include pre-calculated data caching. Pre-calculated data caching may include transformation options for preferred formats such as xml, json and the like. In embodiments, pre-calculated data caching may create schemas. Pre-calculated data caching may create schemas when cloning. In embodiments, pre-calculated data caching may include support for user profile configuration, informing a dashboard about schemas and tools to create a schema and modify it on a dashboard. In embodiments, pre-calculated data caching may include a CRUD editor for tables in a data store. In embodiments, a dstore subsystem may provide securely managed views into SQL databases, a reduction in the number of round trip calls normally required for bulk inserts, compressed queries, and powerful tools for cascading parameterized insertion, in addition to traditional deletion methods. In embodiments, dstore queries may be exposed to feeds. A CRUD editor may be a templated CRUD editor that may use available queries. Queries may be auto-suggested based on tables. Pre-calculated data caching may auto create JavaScript snippets to use in a project for retrieving data, for an application system front-end as well as other systems, such as PHP, website JS, and C# for .NET systems.

In embodiments, a server kit **186** may include a scaling and load balancing capability. The scaling and load balancing capability may request a layer be managed independently of actual instances of the server kit **186**. The scaling and load balancing capability may manage a database layer provided by a third party and maintain a synchronized overall list of files on a third-party system, such as an Amazon S3™ bucket, filename or uploader.

The application system **100** may manage server kit **186** configuration versioning, as well as server kit **186** test and production configurations. In exemplary and non-limiting embodiments, server kit **186** test and production configurations may change views from test to production database systems.

A detailed example of a server kit **186** is provided in further detail below in Appendix B.

The engine **102** may provide an ability to both exploit the benefits of specific devices and provide an abstraction layer above both the hardware and differences among operating systems, including MS Windows, OX, iOS, Android and Linux operating systems.

The editor(s), viewers **144**, and cloud services **142** may provide a full layout, design, deployment system, component framework, JS extension API and debugging system. Added to this the turnkey cloud middleware may prevent

users from having to deal with hundreds of separate tools and libraries traditionally required to build complex applications. This may result in a very large reduction in time, learning, communication and cost to the user.

Unique to an application system **100** described in these exemplary and non-limiting embodiments, projects published using the application system **100** may be delivered in real time to devices with the viewers **144**, including the preview viewer and portal viewer. This may make available totally new use cases, such as daily updates, where compiling and submitting an app store daily would not be feasible using currently available systems.

One aspect of the underlying engine architecture is that the application system **100** may easily be extended to support new hardware and software paradigms. It may also provide a simple model to stub devices that do not support these elements. For example, a smartphone does not have a mouse, but the pointer system could be driven by touch, mouse or gesture in this example. As a result, a user (e.g., developer) may rely on a pointer abstraction. In some scenarios, the pointer abstraction may work everywhere. In other scenarios, the abstraction may only work when required, using the underlying implementation, which may only work on specific devices.

The engine framework may be designed to operate online and offline, while also supporting the transition between states, for example by using caching, syncing and delayed updates.

The engine **102** may provide a choreography layer. A choreography layer may allow custom JavaScript (JS) code to be used to create central elements in the editor, including custom components, custom feeds, and custom actions. In some implementations, JavaScript may be a performance bottleneck, but this choreography approach means JS may be used to request the engine to download files, perform animations, and transform data and images. Because the engine **102** may handle a majority of the processing, the engine **102** may apply all of the platform-specific code, which is able to best exploit each device and manage low level resources, without requiring user involvement.

In embodiments, a domain-specific declarative language of the application system **100** may be important for internal application requirements. The domain-specific declarative language **140** may be used to develop the shared code for the editor, preview and portal applications. These applications may be fully-compiled with LLVM.

In embodiments, the visual editor **108** may be configured to also serialize code for the application system **100**. Serializing content code may allow users to actually edit and load user projects into the engine **102** at runtime. In embodiments, serialization may refer to the process of translating data structures, objects, and/or content code into a format that can be stored (e.g., in a file or memory buffer file) or transmitted (e.g., across a network) and reconstructed later (possibly in a different computer environment). In embodiments, the ability of the domain-specific code to be both a statically compiled language and a dynamic style language capable of being modified at runtime enables the engine **102** to allow users to edit and load user projects at runtime.

In embodiments, the domain-specific language **140** may contain the ‘physical hierarchical’ information on how visual elements are geometrically nested, scaled and rotated, essentially describing a ‘scene tree’ and extending to another unique feature of the language, which is explicit state information. Explicit state information may declare different properties or method overrides, based on different states. Doing this is an example of how an application system **100**

may be able to formalize what current state-of-the-art systems would formalize using IF-style constructs to implement and process conditional logic.

In embodiments, underlying rendering may be performed using a full 3D stack, allowing rendering to be easily pushed into stereo for supporting AR and VR displays.

In embodiments, debugging the state of the engine **102**, its objects and their properties, as well as the details of the JavaScript execution, may be done via the network. This may allow a user of the visual editor **108** to debug sessions in the visual editor **108** or on devices running a viewer.

The networking stack, in conjunction with a server for handshaking, may allow the visual editor **108** to be run in multiuser mode, with multiple users (e.g., developers) contributing live edits to the same projects on the application system **100**. Each edit may be broadcast to all user devices and synchronized across the user devices.

In embodiments, “copy and paste” may utilize the ability to serialize code for the application system **100**. This may provide a user with the ability to simply select and copy an item (e.g., an image, video, text, animation, GUI element, or the like) in the visual editor **108** and paste that content into a different project. A user may also share the clipboard contents with another user, who can then paste the content into their own project. Because the clipboard may contain a comment on the first line containing the version number, it would not corrupt the project the content is being pasted into. The line may also have the project ID, allowing resources like images, that may be required, for example by being specified in a property, may actually be downloaded and added into a target project.

Applications may have the same optical result on any system, as the graphical rendering may be controlled down to low level OpenGL commands and font rasterization. This may allow designers to rely solely on the results of live editing on their computer, even when a smartphone device profile is selected. This unified rendering may provide a shared effects system. This may allow GPU shaders, such as vertex and pixel shaders, to be applied to any object or group of objects in a scene tree. This may allow tasks like realtime parameterized clipping, color correction, applying lighting, and/or transformation/distortions to be executed. As a result, users may rely on the environment to treat all media, GUI elements and even navigation elements like a toolbar, with the same processes. It is noted that in embodiments, the system **100** may implement a graphics API that may support various operating systems. For example, the graphics API may provide DirectX support on Windows and/or Vulkan support for Android, iOS and/or Windows.

In embodiments, the application system **100** may enable users to create content based on functional, gestural and cinematic interactions. Functional content may be functionality delivered with a minimum of features. Gestural content may utilize minimal interfaces, physics and direct manipulation. Cinematic content may utilize playful interaction design, art design, sound design and effects.

Functional content may be the backbone of content created by a user. Implementing functional content may require a balance between designing an experience that delivers the key functionality without weighing a user down in options, learning, terms, menus, dialogs or steps.

Gestural content may focus on providing a great experience the first time, and each time thereafter. Gestural content may be a contrast to tiny atomic graphical user interfaces (GUI’s), which may incorporate many, many small buttons and menus. Gestural content may focus on providing broad simple motions, while being agnostic to the input devices

being used. A gestural content interface may preferably be characterized by low complexity, clarity of purpose, broad ranges of input types and clear feedback.

Achieving a cinematic application may include high quality graphics and presentation elements, providing animation (such as on completion of actions), and/or providing a clear state and flow for the user experience. The purpose of cinematic content may be to enable a user to know what the user is doing and where the user is within a flow of activity, such as following a narrative or following a process through a series of steps to a conclusion.

Content created by the application system **100** may achieve these characteristics when a user can start an application for the first time, understand its purpose and intuitively use it. The created content should feel natural, clearly illustrate actions that the user initiates, and look and sound appealing to an intended user.

As depicted in FIG. 3, the continuum of interaction for natural user interface (NUI) interaction is very large. NUI describes a range of philosophies that center on making the user interface intuitive, so that little or no explicit thought is required for the user to undertake actions. NUI interfaces can be used with any suitable input devices, including, but not limited to, keyboards, mice, touchpads, multi-touch sensors, near touch sensors, voice interfaces, pens, gaze trackers, and 3D input mechanisms. The use of machine learning, machine vision, speech processing, big data queries and natural language processing, are all part of making these inputs seem to be more natural as the application system **100** can increasingly take more types of 'human' input.

The application system **100** may enable a user to access the continuum of NUI interaction and include it in the created content. For example, achieving the goal of gestural software may require the incorporation of broad motions, minimal interfaces, tactile interfaces, ergonomics, performance, state, actions, consistency, rich actions and worlds. Broad motions are important and may mean users should be able to use software by sweeping or touching with a mouse, hand, pen or finger. They should not be required to find and locate very small parts of the screen. Double clicks and right clicks are for shortcuts or accelerators only, not for core functions.

Minimal interfaces may have a small number of elements and may present a clear purpose. Generally, this may mean when designing a layout to be reductive, reducing the number of individual elements may be necessary.

To allow for such applications, which still present very rich content, use of compound elements may be implemented. While these may exist as only one element on a layout, compound elements may contain many homogeneous content items which behave the same. An example of a compound element may be a grid or glass shelf containing many album covers. In terms of interaction with a compound object, these usually support more content than is visible at first glance, often supporting navigation through scrolling, panning, or zooming. This means the individual items within the compound element may be limited to being able to be selected, preventing them from conflicting with any navigation.

Tactile interfaces may provide an important mental connection to the physical world for users. In embodiments, the application system **100** may provide ways to improve the ability of applications being developed with the application system **100** to be more tactile, such as direct manipulation, physics of momentum, friction and elasticity, enveloping and the like. Direct manipulation may provide a user with a

direct, rather than abstract, interface. Examples of direct manipulation may include allowing users to use two fingers to resize and rotate elements, multi-finger crumple gestures, dragging, pointing, clicking, and the like.

Physics of momentum, friction and elasticity may be enabled in the physics engine **152** of the engine **102**, such as to allow elements in an application to slide, bounce, roll, accelerate/decelerate, collide, etc., in the way users are used to seeing in the natural world (including representing apparent three-dimensional movement within a frame of reference despite the 2D nature of a screen).

Enveloping features may be provided, such as modulating an action based on a variable input such as pressure/angle of a pen input or the proximity/area of a finger. This may result in natural effects like line thickness while drawing or the volume of a virtual piano key. This may add significant richness to user interactions.

Ergonomics may also be important to applications being approachable. Similar to the process of designing physical products, the layout, size and interaction of elements/controls may be configured for ergonomic factors. With respect to ergonomic factors, each type of input sensor (including virtual inputs such as voice) that a user (e.g., developer) may choose to enable may have different trade-offs and limitations with the human body and the environment where the interaction occurs in. For example, requiring people to hold their arms up on a large wall mounted touch screen may cause fatigue. Requiring a user of a smartphone to pick out small areas on the screen can be frustrating and putting the save button right next to the new button may increase user frustration. Testing with external users can be very valuable. Testing internally is also important. It is even better when quality assurance, developers, designers and artists of content can be users of the content.

A big challenge in ergonomics is the increasing layout challenges of different resolutions and aspect ratios on smartphones, tablets, notebooks, desktops and consoles. It is a significant task to manage the laying out of content so that it is easily consumed and to make the interface handle these changes gracefully. A snap and reflow system along with content scaling, which can respond to physical display dimensions, is a critical tool. It may allow suitable designer control to make specific adjustments as required for the wide variety of devices now available.

In embodiments, the application system **100** may support performance of 60 fps frame rate and 0 fps idle frame rate. Content created by the application system **100** may enable fast and smooth performance when necessary and reduce performance and device resource consumption when possible.

Input, display, file, audio, network and memory latency are typically present. In embodiments, the application system **100** may be configured to understand and minimize these limitations within the engine **102** as much as possible, then to develop guidelines for app developers where performance gains can be made, such as within a database, image handling and http processing. Performance gains within a database may include use transaction and use in memory mode for inner loops, and streaming images and providing levels of detail within image handling. HTTP processing performance gains may include using asynchronous processing modes and showing users a progress indicator.

Among other features, the declarative language **140** may include capabilities for handling parameters relating to states **130**, choices, and actions. In many applications, state **130** is an important parameter to communicate in some way to

users. For example, in an application for painting, a relevant state **130** might be that a pen is selected with a clear highlight, and the state **130** may also include where the pen is in place. A choice may be presented to a user so that content created by the application system **100** may be required to refocus the user on the foreground and drop out the background. By making state **130**, and the resulting choices or action options clear, intended users may find applications developed by the application system **100** more intuitive and less frustrating.

Actions are often how a user accomplishes tasks within applications. It may be important to show users these actions have been performed with direct feedback. This might include animating photos into a slideshow when a user moves them, having a photo disappear into particles when it is deleted, animation of photos going into a cloud when uploaded and the like. Making actions clear as to their behavior and their success makes users much more comfortable with an application. In embodiments, the aim is to make this action feedback fun rather than annoying and not to slow down the user. The declarative language **140** may be designed to allow a developer to provide users with clear definitions of actions and accompanying states **130**.

Consistency is important for many applications. The application system **100** may make it easier for developers to share metaphors and components across families of applications. The application system **100** may provide a consistent mental model for developers and the intended users of their applications **178**. For example, when a user reaches out and touches screen content created by an application system **100**, the content is preferably configured to act as a user expects it to act.

Rich content is often appealing. The application system **100** may bring content created by or for an enterprise, third party content, and user content to the front and center. Examples may include full screen views of family videos, large thumbnails when browsing through files, movie trailer posters filling the screen, and the like.

Worlds are an important part of the human brain. Users may remember where their street is, where their home is, and where their room is. Thus, in embodiments, the application system **100** may be configured to render virtual 3D spaces where a user may navigate the virtual space. This does not necessarily imply that the application system **100** may render big villages to navigate. In embodiments, 2D planes may still be viewed as the favored interaction model with a 2D screen, and even in a 3D virtual world. In embodiments, the application system **100** may enable 2D workspaces to be laid out in a 3D space allowing transitions to be used to help the user build a mental model of where they are in the overall space. For example, as a user selects a sub folder, a previous folder animates past the camera and the user appears to drop one level deeper into the hierarchy. Then pressing up a level will animate back. This provides users with a strong mental model to understand they are going in and out of a file system.

The engine **102** of the application system **100** may draw on a range of programs, services, stacks, applications, libraries, and the like, collectively referred to herein as the third-party modules **146**. Third-party modules **146** may include various high quality open libraries and/or specialized stacks for specific capabilities that enhance or enable content of applications **178**, such as for scene management features, machine vision features, and other areas. In embodiments, without limitation, open libraries that can be used within, accessed by, or integrated with the application system **100** (such as through application programming inter-

faces, connectors, calls, and the like) may include, but are not limited to, the following sources:

Clipper: Angus Johnson, A generic solution to polygon clipping
 Comptr: Microsoft Corporation
 Exif: A simple ISO C++ library to parse JPEG data
 hmac_sha1: Aaron D. Gifford, security hash
 lodepng: Lode Vandevenne
 md5: Alexander Peslyak
 sha: Aaron D. Gifford
 targa: Kuzma Shapran, TGA encoder/decoder
 tracer: Rene Nyffenegger
 utf8: text encoding, Nemanja Trifunovic
 xxhash: fast hash, Yann Collet
 assimp: 3D model loader, assimp team
 Poly2Tri: convert arbitrary polygons to triangles, google
 box2d: 2D physics, Erin Catto
 bullet: 3D physics, Erwin Coumans
 bzlib: compression, Julian R Seward
 c-ares: async DNS resolver, MIT
 curl: network library, Daniel Stenberg
 freetype: font rendering library
 hss: Hekkus Sound System, licensed
 libjpeg: jpeg group
 json: cpp-json
 litehtml: html5 & css3 parser, Yuri Kobets
 openssl: openssl group
 rapidjson: Milo Yip
 rapidxml: Marcin Kalicinski
 sgitess: tessellator, sgi graphics
 spine runtime: Esoteric Software
 sqlite: sql database engine
 uriparser: Weijia Song & Sebastian Pipping
 zlib: compression, Jean-loup Gailly and Mark Adler
 zxing: 1D and 2D code generation & decoding
 compiler_rt: LLVM compiler runtime
 glew: OpenGL helper
 gmock: testing, Google
 googlemaps: OpenGL mapview, Google
 gson: Java serialization library, Google
 msinttypes: compliant ISO number types, Microsoft
 plcrashreporter: crash reporter, plausible labs
 realsense: realsense libraries, Intel
 Eigen lib: linear algebra: matrices, vectors, numerical solvers
 Boost: ADT's
 Dukluy & Duktape: JavaScript runtime
 Dyncall: Daniel Adler (we use this for dynamic function call dispatch mechanisms, closure implementations and to bridge different programming languages)
 libffi: Foreign Function Interface (call any function specified by a call interface description at run-time), Anthony Green, Red Hat
 llvm: re-targetable compiler
 lua, lpeg, libuv: Lua options as alternative to JavaScript
 dx11effect: helper for directx
 fast_atof: parse float from a string quickly.

In embodiments, the engine **102** may be designed to allow libraries and operating system level capabilities to be added thereto.

FIG. 4 depicts a basic architecture of a project built by a content and application creator **109** of the application system **100**. A project built by the content and application creator **109** may use the capabilities of editor **108**, the publisher **138**, the language **140**, and the engine **102** to create one or more items of content and applications **178**. The project may access various creative assets **402**, such as content of an

enterprise, such as documents, websites, images, audio, video, characters, logos, maps, photographs, animated elements, marks, brands, music, and the like. The project may also include one or more scripts **404**.

The application system **100** may be configured to load, play, and render a wide range of creative assets **402**. Most common formats of creative content (e.g., images, audio, and video content) may be supported by the application system **100**. The application system **100** may also include support for various fonts, 3D content, animation and Uni-
code text, etc. Support of these creative assets **402** may allow the application system **100** to support the creative efforts of designers to create the most rich and interactive applications possible.

The script **404** may be implemented in a language for describing objects and logic **126**. The language may be designed with a straightforward syntax and object-oriented features for lightweight reuse of objects. This may allow the application system **100** to require relatively few keywords, and for the syntax to follow standard patterns. For example, a declare pattern may be: [keyword] [type] [name]. This pattern may be used to declare visible objects, abstract classes, properties or local variables such as:

```
instance image my_image
end
define image my_imageclass
end
property int my_propint=5
int my_localint=5
```

The language **140** may be kept very simple for a user to master, such as with very few keywords. In embodiments, the language **140** may be fewer than thirty, fewer than twenty, fewer than 15, fewer than 12, or fewer than 10 keywords. In a specific example, the language **140** may use eleven core keywords, such as:

```
Instance
Define
If
Loop
Method
Raise
Property
End (ends all blocks)
In (for time based method calls, or assignments)
Tween (used for time based assignments which will be
  animated)
State
```

In embodiments, the language **140** uses one or more subsets, combinations or permutations of keywords selected from among the above-mentioned keywords. An object-oriented syntax may allow for very simple encapsulation of composite objects with custom methods and events to allow for lightweight reuse.

Base objects and composite objects built from base objects may support declarative programming. Support of declarative programming may allow other users, who may be using the GUI, to create functional visual programs by creating instances of a developer's objects, setting properties, and binding logic to events involving the objects.

In embodiments, the language may be strongly typed and may allow coercion.

In embodiments, the engine **102** may be or be accessed by or made part of the editor and runtime infrastructure **104**, to which the language **140** is bound. The engine **102** may support the creation and management of visual objects and simulation, temporal assignments, animation, media and physics.

In embodiments, the engine **102** may include a rich object model and may include the capability of handling a wide range of object types, such as:

```
Image
Video
Database
http
http_server
Sound
3D
Shapes
Text
Web
I/O
Timer
Webcam
```

In embodiments, the application system **100** may provide a consistent user experience across any type of device that enables applications, such as mobile applications and other applications with visual elements, among others. Device types may include iOS™, Windows™, and Android™ devices, as well as Windows™, Mac™ and Linux™ personal computers (PCs).

In embodiments, the engine **102** may include an engine network layer **192**. The engine network layer **192** may include or support various networking protocols and capabilities, such as, without limitation, an HTTP/S layer and support secure socket connections, serial connections, Bluetooth connections, long polling HTTP and the like.

Through the multi-user sync capability **120**, the application system **100** may support a multi-user infrastructure that may allow a developer or editor to edit a scene tree simultaneously with other users of the editor **108** and the editor and runtime infrastructure **104**, yet all rendered simulations will look the same, as they share the same engine **102**. In embodiments, a “scene tree” (also sometimes called a “scene graph”) may refer to a hierarchical map of objects and their relationships, properties, and behaviors in an instantiation. A “visible scene tree” may refer to a representation of objects, and their relationships, properties, and behaviors, in a corresponding scene tree, that are simultaneously visible in a display. An “interactive scene tree” may refer to a representation of objects, and their relationships, properties, and behaviors, in a corresponding scene tree, that are simultaneously available for user interaction in a display.

In embodiments, code may be serialized, such as by a serialization engine **112**. Serializing code may allow a user to create an application in the visual editor **108**, save it, and make changes on the drive of the user, and the change may appear in the visual editor **108**, or a runtime editor **110**. Assets, such as maps, models, scripts, videos and fonts may be synchronized up to a bucket within the cloud associated with a server in a cloud services **142**, such as an S3™ bucket. Modifications may be date stamped against the user who has made them. Modifications may be stored in an undo/commit history of a user, so the application system **100** can rollback changes or deal with versions specific to an individual user.

In embodiments, a frontend may use a real time network system, such as the COMET™ long polling system or a peer-to-peer socket connection, to support synchronization, to tell other computers that a change has happened on the server, which is then pushed down. If not syncing with the server, then the network may send a compressed change down to a local machine. This may allow an application system **100** to work in a propagation mode or a local change

mode. In embodiments, a peer-to-peer socket connection or a Raknet (gaming) layer may support faster, real-time synchronization.

In embodiments, the application system **100** may include an engine UI **106**. The engine UI **106** may be built on top of the engine **102** and may provide an interface to the engine **102**, the editor **108**, the runtime infrastructure **104**, and the like, as well as to individual components thereof. The engine UI **106** and visual editor **108** can, therefore, take advantage of all features of the engine. An application may be created in the visual editor **108** and may be hosted by the engine **102**. In embodiments, the editor UI **106** may also run on the same engine **102** at the same time. This may be supported by introspective and container capabilities of the application system **100**. This capability of the application system **100** may take advantage of an abstraction layer or the abstraction engine **118**, such as for rendering abstractions of visual primitives and for input abstractions of any type of user input.

In embodiments, abstraction may take place over networking and may be an important requirement for simplifying the challenges users may face when creating mobile experiences. In embodiments, the application system **100** may provide simple, object-oriented wrappers for basic lower-level protocols such as http, sockets, serial, MIDI and Bluetooth to address the need for simplifying the challenges developers may face when creating mobile experiences. In embodiments, the application system **100** may also provide higher-level abstractions that may allow users to avoid having to interact with these lower-level protocols. The higher-level abstractions may remove the user from even having to understand which protocols are being used and why.

Examples of the kinds of critical high-level network behavior provided by an application system **100** may include syncing resources and assets; real-time sending and receiving of custom message; enumerating available data feeds and API services; high performance edge caching of data; and compressing data and assets so that networking speed is increased. In embodiments, the engine **102** may embody capabilities of a powerful game engine, which may provide cross-platform abstraction to enabling rendering of highly visual, animated applications, rather than just games. Abstraction and game engine capabilities may be tuned for the rendering of applications using, for example, the visual editor **108**.

An editor and runtime infrastructure **104** may support the visual editor **108** and may provide the ability to have multi-user execution and serialization (rather than compilation) of applications. This capability means that the application system's declarative description of an app/project may be shared and synchronized between different users of the visual editor **108** and the engine **102**. The engine **102** may support not just the reflective and introspective capabilities to make this possible, but the engine **102** may also support suitable object, network and file properties to make synchronization efficient.

The engine user interface **106** may support a UI layer and may provide an ability to create a common set of GUI elements, which may subclass the basic graphical primitives implemented by the engine **102**. These GUI elements may also include all the behavior and uses of an abstraction layer of an engine for allowing handling of various input types. Input types may include touch, stylus, keyboard, mouse, gesture, voice and the like.

The application system **100** may support the declarative programming language **140**, referred to herein in the alter-

native in some cases as the dynamic language or the declarative language. The declarative language **140** may include objects that may form the skeleton of declarative programs. A developer may create instances of the different classes of objects to form a visual basis of a document, an application, or other content. The developer can then set the properties of these objects to adjust the appearance or behavior of the objects in the document, application, or other content. The declarative programming language **140** may include objects, classes, properties and methods. Objects may be discrete bundles of components, often relating to a visual element (a button) or an abstract real world analogy (a customer). A keyword "INSTANCE" may be used to create an instance of a class, otherwise known as an object. A class may define a new type of object. A keyword "DEFINE" may be used to create a sub-class which may be based on an existing class. A property may be an attribute of an object represented by a value, which has one of the types defined for the application system **100**. A method may be one or more actions that may be performed on an object. Such a method may take parameters (like a function) to help describe what the method should do. A method may contain code and may return values.

A class may raise an event, which may trigger a method in an instance of the class. A code may be an ordered list of instructions. A code may be performed (e.g., executed). A code may modify objects or properties and/or may call methods. Objects may be 'nested' into hierarchies. For example, a button may be placed inside a panel. As a result, visual properties such as position and scale may be inherited from the panel. When the panel is moved, the child button may move with it. Objects may include methods. Methods may contain code, which may perform operations. Methods may also bind to events, which may be raised by users. An example event may be a user clicking a mouse. In response to such an event, a method may perform a function tied to the event, as defined by the code.

The declarative programming language **140** may include sub-classing. Sub-classing may include creating a sub-class from a parent class. This is also referred to as a new class 'inheriting' from its parent class. This may be used to create new and more complex classes from the original object model, for example. New programmer-created classes may be used in the language **140** and may define custom properties, methods, and events.

A script **404** of the application system **100** may be made in declarative language. The declarative language **140** may allow a user to define the layout of the hierarchy of objects and their properties that the user wants the engine **102** to create. The declarative language **140** may include an instance. As depicted in FIG. **5**, an instance may make an object of class image, which may make an actual instance of an apple graphic.

An instance may be a nested instance. As depicted in FIG. **6**, a nested instance may make an object of class image, which may make another instance inside it. In the example of FIG. **6**, an object of class image corresponding to a worm is nested in the object definition of a class image corresponding to an apple.

In embodiments, the declarative language may include a "define" function. A define function may allow a user to define a new sub-class of image, which may not create an actual object. FIG. **7B** depicts an example of a define function.

In embodiments, the declarative language **140** may include a "define-and-raise" function. The define and raise function may bind a new class to handle a mouse-down

event and, in turn, raise a custom on-release event when an instance is created from this new class. FIG. 7B depicts an example of a define-and-raise function. The declarative language **140** may also support overriding a base method in a sub-defined class. This allows for compile time checking of a method called in this manner and is equivalent in most cases to using raise. Raise still has the benefit of finding a method when called on an unknown class, but has the disadvantage of needing a runtime type check, so it is marginally slower.

The declarative language **140** may include a logic function **126** (also referred to as logic **126**). A logic function **126** may be a series of instructions (e.g., code) that may each perform basic tasks, which when combined, may create complex behaviors.

The declarative language **140** may include types. A type may apply to the type of an object (e.g., an image), a basic variable (e.g., string, int, real) and the like. Types may be basic types. Several basic types may be available. Basic types may be simpler than full objects and may include:

- Int (int8, uint8, int16, uint16, int32, uint32, int64, uint64)
- Real (Real32, Real64)
- String
- Bool
- Vec2, vec3, Mat4
- Color
- Datetime
- Var

The declarative language **140** may include expressions and assignments. An expression may be a combination of numbers, variables, methods and operators, which evaluate to a value. An assignment may include setting a property or variable to a new value (potentially from an expression).

Expressions and assignments may allow a variable or property to be assigned a value. The value may be an expression as simple as a single number or contain operators like addition, or calls to methods, which return values. FIG. 8 depicts examples of an expression.

The declarative language **140** may include conditions. A condition may be based on a comparison succeeding or failing branch the flow of the program. Conditions may be decisions where code may be branched based on the success or failure of a comparison. If, else, and elseif may be the structures provided. The condition may need to evaluate to nonzero to be true and succeed. Comparison operators may be used to test values. Comparison operators may include: ==(equal), !=(not equal), >(greater), >=(greater or equal), <(less than), <=(less than or equal). FIG. 9A depicts a simple “If block” condition. FIG. 9B depicts an “If/Elseif/Else block” condition.

The declarative language **140** may include loops. A loop may repeat instructions until a condition is met. Loops may allow repeating of a block of logic/code multiple times. A loop may instruct code to run until a condition is met. Loops may be pre-tested and post-tested. FIG. 10A depicts a simple pre-tested loop or “(Loop if)”. FIG. 10B depicts a posted test loop or “(End if)”. FIG. 10C shows an iterator style loop, and FIG. 10D shows looping over a map or list.

The declarative language **140** may include arrays. An array may be a variable that holds more than one value declared by adding square brackets. Arrays may be used with loops and may process lists or maps of data. The application system **100** may support numeric and string indexes. To access the specific item in a condition, expression, or assignment, an index may simply be specified between square brackets. Note that the underlying data structure of the array as a list (such as involving a numeric

index) or map (such as involving a string index) can be selected for capability and performance purposes. Multi-dimensional arrays may also be declared and used by having multiple sets of square brackets. FIG. 11A depicts an example of declaring a basic array. FIG. 11B depicts an example of accessing an array. FIG. 11C depicts an example of using an array in a loop. FIG. 11D depicts an example of declaring a two-dimensional array.

The declarative language **140** may include the capabilities for creating variations **190** and for managing states **130**. A state **130** may provide a state dependent scope (e.g., one can override properties or methods based on a current state **130**). A variation **190** may be a mechanism that allows a developer to override one or more properties, or methods in an object given a state **130**. A variation may be used, for example, to change a style, a layout, a background, a font, displayed text, images, interface elements, and the like.

To ensure that variations **190** may be statically declared, a language feature may be introduced to provide a “conditional scope.” Objects may have multiple pre-declared named states **130**, where in a particular state one or more properties or actions may be overridden. This may allow the application system **100** to handle state-dependent capabilities, such as for localization. It may also be used to handle any condition where a different presentation or behavior may be required or desired. FIG. 12A depicts an example of a variation **190** and state **130**. FIG. 12B depicts how the “language_german” overrides the “width” property and the “caption”, while also adding a new behavior for the actions. FIG. 12B also depicts the revised behavior of the engine **102** in a situation involving a variation **190** or a dependence on state **130**. In this example, the “add_button” object starts with initial properties and methods. Setting the state **130** property of “add_button” to “language_german” applies new properties and methods, without resetting any properties that are not specified as being dependent on the state **130**. States **130** may be added to definitions and instances and around methods.

The variations engine **190** may use variation rules to apply changes to the layout of a page or a master page (which multiple pages can be sub-classed from). The variations engine **190** may enable modifications to control regions and other objects or master objects (which multiple objects can be sub-classed from) contained on these pages. In embodiments, the types of variations that may be possible may include, without limitation, to hide or show an item, to set a property, to request a custom action for this object, or the like. This may be visualized as a checklist by the designer, such as by using the visual editor **108** of the application system **100**, for example. FIG. 13 depicts a list of some example variation rules. In addition to the exemplary rules shown in FIG. 13, variations **190** may be triggered based on the presence of various combinations or permutations of features of the runtime environment or device on which content or applications **178** will run or be rendered (e.g., combinations operating system, region, manufacturer, orientation, pixel density, device and/or language). For example, a variation **190** may be triggered based on the recognition that a mobile app created by the application system **100** is running in an iPhone™ located in Germany, or that an animation is being rendered on a Samsung™ display, etc. The item to which a variation **190** is applied is referred to herein in some cases as the target. The target may be a page, a master page, or the like and may refer to any item contained therein.

States **130** can be established by checking the rules upon startup and after changes to any of these attributes. The rules

may be evaluated in order and each rule may maintain a state **130** of how it has been applied so that when it is applied the first time, the rule matches and is correctly reset when the rule stops matching, for example. The result is that designers may simply create lists of variations (and the states these apply to) to specify the detailed design, layout, or behavioral changes. This list may easily be reviewed and discussed within teams of non-technical staff.

The application system **100** may include a visual code editing environment or the visual editor **108**, as depicted in FIG. 1A. A visual code editing environment may allow a developer to work at a high level while an underlying engine **102** has full control of CPU/GPU utilization to the lowest level for best performance.

In embodiments, the visual editor **108** may allow the application system **100** to produce detailed analysis at the CPU level and analyze assembly level instructions to maximize thermal performance of applications (at the level of individual instructions), enabling full control of an entire stack of resources down to the lowest level, including CPU utilization and instructions generated by a compiler.

The application system **100** fundamentally may change the process of delivering information to endpoint devices and create an engaging interactive experience. The visual editor **108** may allow multiple groups to access simultaneous layout and structure information that is packaged into the container **174** for publication. This may be done in minutes for simple projects, and in no more than a few days for projects that draw extensively from external information feeds or that use complex graphics. The container **174** may be instantly published to a computing cloud, such as an AWS™ computing cloud, where it may then be immediately accessed by authorized users, on any device, with an identical experience, or with an experience that reflects desired variations **190** that are particular to the situation of the user. The container **174** may also be published to a private or public app store, in which case it may be subject to the standard approval process and other requirements of the app store.

The application system **100** may include a gaming engine **119** and other capabilities for handling machine code across platforms. The gaming engine **119** may be implemented with a low-level device (e.g., a chip, ASIC, FPGA, or the like), performance-based approach on the machine side and a gaming engine approach for visual elements, such as using the abstraction engine **118**, to marry visual behavior to a machine environment. The gaming engine **119** may be highly optimized for applications which have a large number of GUI elements. These elements may be 2D, 3D, or a mixture of both. The gaming engine **119** may also be able to manage transitions and animations very efficiently. The gaming engine **119** may provide capabilities that may have been ‘hard-coded’, which with the gaming engine **119** is parameterized, such that it can be varied for execution as desired according to the capabilities, state, or the like of an execution environment for content or applications **178** created using the application system **100**. This allows content or applications **178** in a project to be serialized, without needing to be compiled when it is distributed to the viewer/portal **144** or the editor **108**.

The application system **100** may include a platform for code development. The application system **100** for code development may include a plug-in system **121**, such as a JavaScript plug-in system, the content and application creator **109** (including the editor **108**), the script layer **134**, and the engine **102**. The plug-in system **121** may orchestrate components, feed events, and trigger behavior. The content

and application creator **109**, including the editor **108**, may integrate with or connect to the script layer **134**. The script layer **134** may be bound to an assembly language **123** that may be bound to the high-performance engine **102**. The LLVM compiler **136** may compile code to communicate with the engine **102**.

The engine **102** may use JavaScript to trigger behavior or respond to an event. The editor **108** may have rich function components. The script layer **134**, such as using JavaScript, may orchestrate those components while rendering at native speed. New behavior may be rendered in an application without changing a binary; that is, what was compiled from the editor **108** and downloaded, may remain unchanged, while only the script, e.g., JavaScript, changes.

The engine **102** may include low level engine features like the animation and simulation. In an example, the following JavaScript call is equivalent to the internal application system **100** “tween” statement:

```
setPropertyInTween(apple, "alpha", 1.0, 1000, 5)
```

The foregoing “tween” statement is equivalent to the statement:

```
apple.alpha=1.0 in 1000 tween ease_both
so the component apple has its alpha property changed to 1.0
over 1000 milliseconds using the ease_both tween.
```

The application system **100** may provide an extension API **125**, such as for allowing actions to be extended, such as using a scripting language like JavaScript. This may allow new classes of components and data feeds to be created, such as in JavaScript.

The following example depicts how JavaScript is used for creating one of the basic components and then adjusting its properties. The result of this example is the full 60 fps hardware accelerated performance of the base component, as well as the custom behavior the following JavaScript provides:

```
registerComponent("toggleButton", "", "toggleButton", "", "toggleButton_
init", "", "toggleButton_refresh",
"imagepicker:default_image,imagepicker:down_image,imagepicker:
disabled_image,
imagepicker:toggled_default_image,imagepicker:toggled_down_image,
imagepicker:
toggled_disabled_image,bool:show
text:0,string:default_text:on,string:toggled_text:off,bool:toggle_state:0", "")
function toggleButton_init(width, height){
45     var toggleButton = createComponent(self, "button", 0, 0, 100;
        100);
        var overlay = createComponent(self, "image", 0, 0, 100, 100);
        bindEvent(overlay,"on_press","togglePress");
    }
function togglePress( ){
50     var button = getComponentChild(self,0);
        var toggleState = getData(self, "toggle_state");
        if(!toggleState){
            setProperty(button,"default_filename",getData(self,
"toggled_default_image"));
            setProperty(button,"down_filename",getData(self, "toggled_down_
image"));
55             setProperty(button,"disabled_filename",getData(self,
"toggled_disabled_image"));
            setData(self,"toggle_state", 1);
        }
        else{
60             setProperty(button,"default_filename",getData(self,
"default_image"));
            setProperty(button,"down_filename",getData(self, "down_image"));
            setProperty(button,"disable_ filename",getData(self,
"disabled_image"));
            setData(self,"toggle_state", 0);
65     }
    }
```

In addition, specific capabilities of the engine **102**, such as animation and simulation, are exposed to the JavaScript API just as they are exposed to the declarative language **140** of the application system **100**. Hence a statement like `setPropertyInTween (apple, "alpha", 1.0, 1000, 5)` will execute very efficiently, as it is only called once and, subsequently, the hardware accelerated animation system inside the engine **102** takes over and performs all the animation and rendering of the transition of this single (or many) property change, including multiple tweens inserted in the timeline for this (or many) other objects simultaneously.

In embodiments, the application system **100** may allow serialization of the content from an editor, which may allow a user to use the full power of the underlying engine **102** and its gaming engine features (without the need to compile anything on the client).

In embodiments, the engine **102** may allow a user to better abstract from the underlying operating system. Traditional, native development environments bind to the native GUI library provided by Apple, Google, Microsoft, etc. This results in very different behavior, look and feel, and programming models on each native device. The application system **100** may avoid all of that by creating all GUI elements out of visual primitives fully controlled by the engine **102**. This makes the pixel level designer control and cross-platform performance very effective in the application system **100**. For example, a toggle switch in iOS™ looks like a white circle in a green oval, which slides left/right. Yet in Android™, it is a box within a rectangle, which slides up and down. In Windows™, this is often a 'check' icon in a box which is present or not. In application system **100**, the visual designer may pick and customize their own pixel perfect look, which works on all devices.

In embodiments, the visual editor **108** may use the same core engine **102** of the application system **100** for editing, previewing, and running at runtime. This may ensure that there is no difference and that the experience for the user is completely WYSIWYG ("what you see is what you get"), which may reduce testing.

In embodiments, the visual editor **108** may use the same engine **102** while editing and while running. This may be powered by the LLVM compilation technology of the engine **102** that may allow a user to pre-compile the editor code and the engine editor and runtime infrastructure **104** code using a shared code base which then is bound to the engine **102**. An example implementation of an LLVM compiler is provided in Appendix A.

In embodiments, the application system **100** may share the same core code shared among editing, previewing and runtime environments. This may allow the application system **100** to draw on the same code base for the hosting and bootstrapping of script applications created by the end user, which may not need compilation with LLVM.

In embodiments, the visual editor **108** may allow real-time, multi-user, simultaneous development by allowing a shared understanding and shared simulation of what is being created. In embodiments, the visual editor **108** may include an optimistic locking system. In an optimistic locking system, the last commit may win. An optimistic locking system may restrict, so that only a master user can do some tasks, such as destructive editing when someone is working on it.

Multi-user development may include a parallel process between users. Optimistic locking may allow the application system **100** to make almost everything multi-user. Optimistic locking may be used to manage updates to a central representation of the current project/application. Importantly, this may be able to interact in a transactional model

with all parts of the system so that if it is an asset in the folder, an image or change to a JavaScript file, for example, a new component being added to a page, or a change to a property on an existing component. These can all be kept in-synch between users. Importantly, if a user joins the session late or drops off the connection periodically, the user may get caught up with a list of all transactions that have occurred.

When one of the users manually chooses to save (e.g., publish to the viewer), this may act as a sync point and may collapse the delta/transaction log so it does not become unwieldy.

In embodiments, the engine **102** of the application system **100** may be designed to support the visual editor **108**, which is both an application itself and can inspect and edit another running application that a user is working on.

In embodiments, the visual editor **108** may include a JavaScript API. A JavaScript API may be designed to enable very light configuration work, leaving the engine **102** of the application system **100** to do the bulk of the processing. This may prevent runtime JavaScript from becoming a bottleneck, and instead may take full advantage of a hardware-accelerated engine.

In embodiments, a further network layer via cloud services **142** may provide real time synchronization of assets and changes to the application being edited between multiple users.

In embodiments, the engine **102** may utilize a LLVM compiler **136**, either integrated with or as part of the engine **102**, to act in a just-in-time compilation mode for developers, or without LLVM where it may simply be part of the toolchain to generate a finally compiled application.

The visual editor **108** and viewer/portal **144** both may use the engine **102** as a fully compiled application. This may allow applications developed using the application system **100** to run on systems like tablets or phones, which may not allow runtime compilation.

Both the editor **108** and the viewer/portal **144** may be created using script **404** made and handled by the script layer **134**, which may be compiled using the LLVM compiler on a build wall (cloud hosted CPU's with specific hardware, OS and toolchain configurations dedicated to compiling OS specific executables), which may then be linked with the engine **102** to create final static binaries, which can be installed on devices. A user may synch one of their projects with an editor/viewer/portal **144** and the following things may happen. First, the assets (images, scripts, sounds etc.) may be copied into a cache folder specific to this project. Second, the project file (in compressed script **404**) may be copied into the same cache folder. Third, the engine **102** may parse this high-speed file, thereby creating all the required components and storing lists of required built-in engine actions, or links, to JavaScript actions. Fourth, content or application **178** may now be used interactively; it may be effectively hosted and bootstrapped by the editor/viewer/portal **144**.

The editor **108** and other elements of the editor and runtime infrastructure **104** (such as for preview and portal features) may be written in the declarative language **140** and may be compiled with LLVM **136** and bound to a C++ engine for maximum native performance.

As referenced elsewhere in this specification, the application system **100** may include the declarative language **140** (rather than just using an API). The declarative language **140** may support LLVM as the main compiler. The declarative language **140** may include an LLVM front-end. The use of

LLVM in the application system **100** may increase the efficiency and speed of an LLVM front-end.

The same language **140** may be used to build an editor and an application. It may also be the file format/language which is interpreted, from which the runtime behavior is edited and transmitted and then executed by the engine **102** as a simulation.

The declarative language **140** may be used to describe not just the objects, but their relationship within a scene tree. This may include a hierarchy and nesting of objects (and their properties and methods). This may allow building the full detail of the page layout. Properties may also be inherited so that a form that contains a button, when rotated or scaled, will appropriately rotate and scale its children (in this case the button). The declarative language may be interpreted and executed by engine simulation. FIG. **14** depicts the declarative language scene tree description **124**.

In embodiments, the ability to express both the logic layer and layout/presentation layer within the same declarative language **140** is provided, which may reduce complexity and make separation of logic and presentation optional, rather than requiring developers to handle them separately, which consumes time and can result in errors or unintended effects of the logic upon the presentation layer.

The same domain-specific declarative language **140** may be used to create the editor **108** and runtimes (such as runtimes for a preview of content or application behavior and runtimes that are accessed by users, such as through the portal **144**). That same domain-specific language **140** may also be used to create the content and applications **178**. This may utilize the introspection and container capabilities in the way that the language **140** is implemented by the engine **102**.

The file format **128** may be where a script **404** of the application system **100** may be serialized or de-serialized, such as without the need for compiling. This may be how an app is 'loaded' inside the host environment, such as for preview, in the portal **144**, or in a host environment for editing.

Tokens may be turned into bytecodes and literals (e.g., strings, numbers, and/or object names), and may be stored following a bytecode and length as appropriate. This serialization may be designed to be smaller to transmit and faster to encode/decode without the need for parsing.

A restriction applied to the serialized applications is the logic inside objects. Events may be limited to a list of parameterized methods, to simplify the task of specifying a workflow for the user. With states, users may also have access to the equivalent to conditional logic.

The declarative language may be designed to be extended with new object classes with methods, properties and events which may expose cross-platform device features. This may allow very simple expressions. The declarative language may grant access to logic **126**, databases, animation, 3D, and documents. Language features may have animation, time and multiple states for an object. The declarative language may have a scene underneath, where all properties may be synchronized and run as a simulation.

In embodiments, objects may inherit properties. Examples of objects and properties may include buttons on forms, rotating forms, making elements transparent, and the like. Unlimited named states may be added to an object or object class. Each state **130** may encapsulate properties or methods. This may make it possible to create a variations system. A state **130** may be like conditional logic.

FIG. **15A** depicts a button in German specified using conditional logic. In this example, conditional code may end

up anywhere in a program making it hard to find and debug. The compiler also may not be sure if the code inside the condition will ever be executed as the condition has variables in it. FIG. **15B** also depicts a button in German specified using the declarative language. This may all be declarative and checkable at parse time. The states may be inside the object and may be fully resolvable. A user may declare a list of valid states that the object may be tagged with at the top of a program

The declarative language **140** may include a domain-specific language. A domain-specific language may make it quick and easy to create visual experiences and may assume a user is going to use the framework of the application system **100**.

The declarative language may be extremely fast to compile in a single pass. For example, an entire application using the LLVM compiler may compile in seconds. The declarative language may have the features of a strongly typed declarative language **140**; it may, however, be fully statically compiled for maximum performance. In the declarative language, instead of writing C++ code to create the editor, the entire editor may be significantly smaller than the engine.

In embodiments, the application system **100** may publish apps through a private portal **144** and manage the ability for users to edit applications on the fly. The application system **100** may allow a user to deploy an app without having to go through an app store. A user may compile to the app store. A user may also publish natively to other platforms that anyone owns. Other platforms may include Windows, Android, iOS, Linux (Ubuntu), for digital signage content, for example, and the like. This may allow a user to publish to different platforms seamlessly. This may eliminate the need to put internal facing apps into an app-store or other application vending service. This may allow for development of company-specific portals, where customers may push apps into their own portal, allowing a company to serve content-driven experiences to customers.

In embodiments, the declarative language **140** may support multiple roles. The declarative language **140** may be a strongly typed declarative language for building binary applications, via the LLVM compiler **136**, and for high-speed serialization for the viewer/portal/editor **144**. As depicted in FIG. **1A** the declarative language **140** may include a scene tree description **124**, logic **126**, file format **128**, state information **130**, modules, domain-specific script **134**, LLVM compiler **136**, and publisher **138**. The LLVM compiler **136** may compile with LLVM to final binary code, or publish through a viewer/portal **144** to de-serialize content into a host application such as a portal (or for that matter the editor).

As depicted in FIG. **16**, a scene tree description **124** may support an SGML style nesting of elements. Importantly, the elements may inherit visual and transformation properties from their parents. This is important as moving a panel in 3D space may move all the buttons and text elements pasted on a panel with it. This covers scaling, moving, rotation and changes in visual properties like transparency.

Logic **126** may be placed inside a method. The only distinction between standard methods and event handlers may be when methods using reserved names like `on_click` may receive the appropriate event if it is triggered on that object. Importantly logic **126** may be inherited and separated from presentation as required. The code may natively understand other components and their relationship to each other. This means that in code, a user may explicitly refer to children and their properties using, for example, a “.”

(period) to separate them. For example, clicking on a tree may perform an operation on another object using the dot syntax, each dot may then be another level lower.

As depicted in FIG. 17, at the global level, `tree.apple1.x=100` is valid. At the level of the tree `.apple1.x=100`. At the level of the `apple1` then `.x=100` is valid. It may also be possible to step backward using `.parent()`.

File format 128 may allow the application system 100 languages to perform multiple duties. A sophisticated LLVM compiler 136 frontend may lex, parse, and compile the logic 126. The LLVM compiler 136 may also be designed to be serialized in text or binary format for high performance when used as a file format 128, so as to support publishing through a viewer/portal 144. A file format 128 may have a few specific restrictions like turning logic 126 into lists of parameterized method calls, but largely may be kept the same in terms of sub-classing, nesting, properties and state. The ability to be used as both a compiled and serialized description is part of the language's design and novelty.

The state 130 may be used for various purposes in the application system 100, including to drive the variation system. A state block may provide a scope where any properties or methods specified may override those at the default level. This may allow conditional logic 126 to be statically checked. The state block may check the Boolean state of a flag. This flag may be system derived (`orientation_portrait`) or user defined (`user_platinum_member`).

As depicted in FIG. 18, the example depicted in FIG. 17 is changed to show how statically declared states may now be used to determine the unpicked and picked visualization of the apple. Then the `on_click` method will set the apple to be picked, causing those affected properties to be updated.

Modules 132 may be a set of standard language features which may be found in most languages, e.g., objects, events, properties and sub classing, which may have specific engine level support for both abstracting OS level features implemented in the C++ layer and extending objects with synthetic JavaScript sub classing. Modules 132 may make it extremely quick to add support for new abstractions implemented at the lowest engine level or by third parties as synthetic sub-classes of one or more base classes using the JavaScript API. In embodiments, modules 132 may include novel elements of a python code generation layer, which may be used at the C++ end and more about the JavaScript API harness for `registerComponent` and how it may interact with the engine 102 and its ability to dynamically manage new classes.

A domain-specific script 404 of the application system 100 may be provided within a domain-specific language in the application system 100. The domain-specific script 404 may have multiple special properties. Multiple special properties may include:

- 1) Standard patterns for structures [keyword] [type] [name];
- 2) Simplified keywords;
- 3) Statically declared hierarchical structure like SGML;
- 4) Statically declared sub classing;
- 5) Statically declared conditional logic (state);
- 6) Statically declared animation and timing (In keyword);
- 7) Statically declared access to the components/objects using the dot syntax;
- 8) Serializable to a binary format;
- 9) Single pass lexing and parsing (no backtracking); and/or
- 10) High speed compiling.

An LLVM compiler 136 may cover the way a domain-specific application system 100 language may use a custom

frontend to lex and parse, then may compile a script 404, which may define the GUI and behavior for both the editor/preview and portal from a shared source code base.

There may then be a set of different tool chains for each OS, which may link binary code with an engine, such as the C++ engine, and with OS-specific wrapper functions, as well as with platform-specific bootstrap code (e.g., for C++, Java, Objective C, etc.). This may allow a user of the application system 100 to build a fully compiled native application powered by the engine 102, the GUI and behavior defined in the application system 100 script 404, and the CPU/GPU architecture-specific tweaks present in LLVM backend.

Publishing through a viewer/editor may focus on how a host application works rather than the actual serialization and the benefits that may be provided to a user of the application system 100 (effectively code free development). A host application may be built in the domain-specific script 404 (such as within the editor/viewer/portal 144). This may be able to serialize/de-serialize all the objects/components from part of a tree structure that may make up the scene. This may be how a 'sub-application' is loaded and saved. It is noted that there may be several mechanisms that may make this process work, such as with the dynamic components, methods/actions, and events driven by an API such as a JavaScript API. In embodiments, the application system 100 may include the ability to enter an "edit mode," where component events may be intercepted, so as to allow components to be selected and highlighted. These reflective qualities may be built into an engine object model, and the domain-specific script 404 may exploit these reflective qualities so that once the selection event is intercepted, the positions of components can be interrogated, a highlight displayed around the component, and the properties of the components can be interrogated and displayed in the properties panel and made editable by the user.

It is noted that there may be no compiling going on in the preview/portal. Simply the components may be re-established inside a 'sub-application' by de-serializing them, and the action lists may be bound to the events. This may mean the engine 102 may now run everything at full native performance with the existing binary functionality that the engine 102 provides.

A JavaScript API may add a 'choreography' layer where new functions may be deployed. This API may be designed to offload as much ongoing computation to the engine 102 as possible, keeping the JavaScript layer setting up behaviors or connecting processes like network calls and data processing.

In addition, the application system 100 may include another system that may manage the undo/redo and multi-user systems. This system may be a set of virtual properties and methods. Virtual properties and methods may be shadows of real properties and methods the component has and may not affect rendering activities of the engines. However, these virtual properties may hold the current serialization state and may allow the application system 100 to manage just saving non-default values, resolving changes in transactions during undo/redo or multi-user operations, and resetting components back to their virtual state after a user has tested an app in the editor.

The application system 100 may perform simulations. An application created by the application system 100 may create new objects, animate things, change properties, and the like, during run mode. Being able to return an application back to its editing state using virtual properties may be useful in the application development process.

In embodiments, the application system **100** may create visual avatars (or avatars). Avatars may be ‘digital agents,’ such as ones that assist in answering questions, ones that guide users through a process or workflow, or the like. Avatars may be rendered by an avatar rendering engine **148**. Avatars may interact with a user of the application system **100**, including the engine **102**, by various mechanisms for avatar interactions. Avatar interactions may be generated and managed by an avatar interaction engine **148**. Avatars may include conversational user interfaces (UI’s) where users may productively interact with 3D-rendered visual ‘digital agents’ with human attributes, such as a head which can change angle, an animated mouth and lips to represent human-like speech, eyes that can look and blink, and facial poses that can convey emotions such as surprise, frustration, thinking or happiness.

Current state of the art ‘headless’ systems may provide a poor user experience, by requiring a user to remember many options and phrases to type or speak in performing a wide variety of potential tasks, each of which may have multiple non-explicit stages or requirements. As depicted in FIG. **19**, an avatar and associated visual 3D user interface (UI) **152** may provide a user of an application system with a visual presence and richer input so an application system may take advantage of verbal commands and nonverbal cues provided by human interaction. This may allow an application system to provide a much more useful and productive interface, within an overall 3D UI **152**, with capabilities like presenting options, showing context and conveying progress. FIGS. **19-22** depict examples of how an avatar may be used in an application system, in exemplary and non-limiting embodiments.

As depicted in FIG. **19**, characters may represent the appropriate context of a scenario, for example, adult, child, genders, roles & costumes. Characters may be displayed as avatars by the engine **102**. FIG. **19** depicts an avatar in a 3D user interface **152**.

FIG. **20** depicts how morph targets, such as those utilizing different meshes, may be blended together by an avatar rendering engine **148**, such as using a GPU. In this way, the avatar rendering engine **148** may produce accurate speech animation and overlaying emotion, for example, smiles/frowns, specific actions, eyebrow raises and the like.

FIG. **21** depicts an example of a character interacting in relation to controlling the entertainment console on a car using a backend system that can handle queries directed to the character (such as by conversation), search for entertainment content based on the queries, and respond by offering to play entertainment, or the like. In embodiments, user interaction is managed by a backend artificial intelligence system, such as the IBM Watson™ backend.

FIG. **22** depicts gaze and direction being used to convey that the character is looking at the options and this is what the user should be interacting with.

The engine **102** may include avatar capabilities. Avatar capabilities may include face rendering, voice synthesis, procedural animation, emotion and the like. Face rendering may be performed by the avatar rendering engine **148** of the engine **102**. These capabilities act as a coordination layer above the skeletal animation and morph target systems which most modern 3D game engines support.

An avatar interaction generated by an avatar interaction engine **148** may be an avatar speech interaction. Avatar capabilities may use a voice font and a voice synthesis process on a local device to generate the speech interaction on device to reduce latency.

Avatar capabilities may extract phonemes generated during a text to speech process, for example by mapping each viseme/phoneme to a combination of different percentages of mouth positions, also known as blend shapes, on a 3D model of a face, then animate between these multiple blends in time with the phoneme changes in the generated speech to give an appearance of a talking character.

Avatar interactions generated by an avatar interaction engine **148** may be blended shapes for basic anatomic processes like blinking, emotions and large scale motion like eyebrows, cheeks and the like.

Avatar capabilities may provide a procedural animation system and state machines to drive basic anatomic functions such as head movement, blinking, eyes looking towards the user (via head tracking) when it is waiting for a verbal response from the user, basic semi-random facial motion and the like.

Avatar capabilities may provide a system for keeping an emotional model for an avatar so that based on the conversation the character can react.

Avatar capabilities may include conversational interfaces and natural language endpoints, referred to as conversational interfaces in this specification. Conversational interfaces may support cross API conversational interface processing including a basic command processor.

Conversational interfaces may support cross API conversational interface processing. Cross API conversational interface processing may include capturing audio on a device and sending small chunks of this voice data to a cloud service **142**.

Cross API conversational interface processing may also support receiving back text as a fully processed response by an AI service or simply the raw text after speech to text is completed for further processing.

Conversational interfaces may supply basic NLP (natural language processing) capabilities if required to determine parts of speech and provide basic flexible processing of command trees at the application system **100** level from typed text or text which may have been extracted and potentially partially processed via a cloud based voice analysis process **242** and **243**.

Conversational interfaces may provide a system of attaching actions to parameterized commands from command trees, for example the command “please search for trees” results in action=“search” parameter=“trees”.

Conversational interfaces may provide the ability for the character to turn and look at user interface elements or information that may be part of the conversation/application. User interface elements may be elements of the engine user interface **106**. This is an important human queue which may inform a user when and where they should be paying attention, for example when the avatar is describing something and presenting choices.

An avatar combines a simplified user audio interface with the power and density of information provided with visuals. The application system **100** may include an avatar interaction engine **148** and an avatar rendering engine **148** to capture input, process and present avatars and content.

This presentation of workflow, options and content (images or models) based on the user interactions managed by the avatar interaction engine **148** and state of the current conversation are important—this allows the user to understand where they are in the workflow and what they can do. This simplifies the types of interaction and requirement for users to remember the current state themselves as an audio or text only system does.

Essentially, the avatar interaction engine **148** allows developers to make the avatar appear much more intelligent, responsive and useful to the end user by improving their understanding of what the ‘avatar is thinking’ (by explicitly showing the user) while providing clear commands so that the user can interact by speaking without ambiguity and frustration by speaking these commands rather than expecting the avatar to be able to carry out a fully human like conversation and all the nuances and sophistication that this entails. These commands map directly to the logic **126** defined in the visual editor **108** of the engine **102**.

Avatar interactions, as managed by an avatar interaction engine **148**, may include voice recognition, gaze detection, face analysis, facial expressions, hand gestures, Halo fingertip detection (as discussed later in this disclosure), phoneme generation, phoneme mapping per character, natural language processing and the like.

For a user of an application system **100**, interaction with agents allowing multi-modal input may be critical. For example, a user may read out an option shown on a display, but just like another human, a display may see where a user is looking. An agent should be able to use a front facing camera where possible, to track the fixated eye gaze of a user, and to allow users to use voice commands to trigger an option they are staring at. If a display is touch capable, then a user should also just be able to touch the option they want. High performance eye trackers which can manage saccades and other subtleties of eye movement will also be beneficial, but are not a requirement.

Providing a range of useful interaction choices may allow a user to be more productive and may make an experience more natural.

An application system **100** may include a 3D camera for detecting the eye direction based on the shape of the eyeball including cornea (which is not spherical). An application system **100** may also use basic 2D image processing for iris and pupil ‘glint’ extraction based on light bouncing off an eye from a structured light IR element in the 3D camera (captured using the raw 2D IR view prior to 3D processing from the 3D camera).

Avatar capabilities may include facial analysis. Facial analysis of a user may be used to not just extract gaze tracking, but with the details that may be extracted from a 3D depth camera such as an Intel RealSense 3D camera, an application system **100** may find the jawline, mouth, and nose. This may enable an application system **100** to both determine if a user is facing the computer and if a user is moving their mouth. In embodiments, the application system **100** may further analyze an audio input to alter the detection threshold for audio commands in noisy environment. In embodiments, the application system **100** may further perform facial analysis on the output of a camera (e.g., 3D depth camera) to extract an emotional state at a statistical level.

Avatar capabilities may include recognizing facial expressions. In contrast, characters, expressing emotions like happiness or sadness (on a scale) or performing specific actions like laughing or frowning, through facial expressions, may be very useful. This kind of emotional interaction, which may be conveyed by facial expressions created by an animator or from motion capture, may be easily read by a user. This may greatly improve the verisimilitude and help convey useful or entertaining information to a user which only a character can.

Avatar capabilities may include recognizing hand gestures and pointing. Hand gestures and pointing may provide significant useful information. This may be referred to as

“near touch”. This may be similar to a mouse hover state and may allow a user to express interest without actually touching a surface.

Avatar capabilities may extract not just the fingertip using a machine vision library, for example, but the vector of the direction it is pointing. Using this, vector avatar capabilities may extract the position where a user is pointing relative to a camera and hence a screen (where the geometry in relationship to a camera may be known). This is very powerful as a user may point and an avatar may display more information, for example equivalent to hover tips. The user can then confirm selection by touching, speaking or tapping in mid-air, for example using fingertip tracking.

Avatar capabilities may include phoneme waveform generation. An application system **100** may use a version of festival from CMU for the text to speech process. In embodiments, the application system **100** may utilize a viseme list and voice fonts to determine the types of gender, accent and languages spoken. In embodiments, avatar capabilities may extract the unique mouth poses (phonemes) to be used by the animation system to configure the morph blending aligned with the right timing for the audio layer, for example within 33 ms accuracy. The combination of the generated speech and this facial animation as visual feedback may make it easier for people to understand the speech of a character in noisy environments.

Avatar capabilities may include a phoneme mapping system. A phoneme mapping system may allow unique character designs with non-human faces to be mapped to speech and emotional reactions. An animator may be allowed to specify ratios of up to five blends for each phoneme, for example “wh”, action, for example “blink”, or emotional response, for example “happy” overlay.

Character poses and design may be very flexible and may allow realistic humans, cartoon humans and anthropomorphic object or animals to be created.

As discussed elsewhere in this specification, an application system **100** may include a speech tagging module, command trees, a basic ontology loader, and/or a state management system. These capabilities may enable the application system **100** to perform basic avatar interactions without the need for an online backend. Users of the application system **100** may also need to enter plain text to drive interaction without an operating system level, or third-party, voice recognition system. As such, these tasks may require a natural language processing (NLP) capability.

An application system may include an enhanced NLP capability, which may take advantage of extra information an application system has over text/voice only systems, for example, an application system may have visible commands and visible options which may make a massive difference in removing ambiguity from user interaction and user commands. This may make the user experience, simplicity and discoverability much more effective. For example, the current state of the art systems has very low, for example, 2%, attach rates due to the poor ability to remember the commands to initiate a skill, or the syntax required to make the skills work.

In embodiments, the application system **100** may include a three-dimensional (3D) user interface (UI). A 3D UI may include a mixed 2D/3D scene tree layer. To support applications which mix 2D and 3D elements in the same application, a novel hybrid scene tree system may be used allowing many 2D and 3D elements to be efficiently rendered, composited and interacted within the same visual scene. This may use both a traditional parent and child list for holding the primary relationships of a hierarchy and

managing the propagation of the inheritance of the properties and matrices down a tree. In embodiments, there may also be a unique combination of tree structures for hit detection and scene management for visibility as well as a connection to an alternative tree structure to manage the overall 3D scene helping with 3D picking and 3D occlusion.

In embodiments, a 3D UI may include a library of machine vision primitives. A library of machine vision primitives may be configured from JavaScript, making leading edge capabilities accessible to scripting. Some elements like ghost hands, described later in this specification, 2D/3D barcode detection, fiducial detection and more may be made accessible directly without any code using a library of machine vision primitives. Machine vision primitives may include blob detectors, optical flow, filters, fingertip detectors, and face landmark tracking. A blob detector may be a 3D blob detector. A blob detector may be an RGBZ blob detector and may include color and the distance to the screen. A blob detector may include algorithms to disambiguate gradients, establish angles, and isolate blobs to measure direction, size and movement to sub pixel accuracy.

In embodiments, a 3D UI may extend a rendering engine for stereo and 3D input for VR/AR. Because of the nature of even 2D applications and all the 2D UI existing inside a 3D engine, applications when rendered in stereo and using 3D input like hands or controllers, may work with AR and VR systems.

In embodiments, a visual editor **108** may include a collaborative 3D workspace **300** for users of augmented reality (AR) devices. FIG. 2D depicts a collaborative 3D workspace **300** according to exemplary and non-limiting embodiments. AR devices may be phone-based AR devices, headset-based AR devices and the like. The application system **100** may allow a user to both use 3D interactive experiences, as well as create them in AR using the visual editor **108**. A collaborative 3D workspace may include a fused 3D scene for multiple users. For example, all users of a collaborative 3D workspace may have a shared virtual environment that may map onto the real world. As a user moves around the real world, the virtual 3D may be perceived as fixed at the same location, for example on a table or against a wall. In addition to the shared 3D elements, each user may also be able to see their own user GUI elements and digital avatar rendered from the perspective of the user.

In embodiments, a collaborative 3D workspace may include 2D planes, which in cases may provide a very useful interaction model. Multiple 2D workspaces may be laid out into real 3D environments. Many data types like documents, videos, illustrations and even 'frames' for 3D models may then be placed, organized and edited in the real world via these augmented 2D planes placed into the real world.

Users who may be viewing the augmented world via a mono viewport, such as by looking at a smartphone screen that renders virtual elements over the camera view, for example, may have some specific limitations. While the combination of optical flow, feature tracking, and inertial sensors utilized by augmented reality programs like ARkit™ or Google Tango™, for example, may provide the ability to treat a smartphone screen as a convincing window into the augmented scene, it may be impractical to hold the phone up for any length of time.

This may also apply to stereo headset viewing, which may require cases where a user needs to get closer to a virtual object. As a result, an aspect of a collaborative 3D workspace UI framework includes a zoom UI. A zoom UI allows a user to select a data type and zoom into a mode where only the selected item is locked to the viewport of a user for

viewing and editing. An aspect of a zoom UI is providing an animated transition for a data plane as it zooms into a user viewport or back out to its world position. This may allow a user to keep an overall mental model of an environment, even as the user switches between items in zoom UI mode.

Additionally, in embodiments, it may be critical for stereo views that a zoom UI user viewport is aligned on the zero plane, indicated by zero convergence between eyes, meaning there are no issues with eye strain or eye dominance. This may also enhance a number of low level GPU optimizations, making it possible to reduce thermal overhead on portable AR and VR stereo devices.

In embodiments, a collaborative 3D workspace **300** may include three core elements. The three core elements may be able to be used individually and also benefit from the collaboration that an underlying application system engine **102** and network layer make possible.

Three core elements for a collaborative 3D workspace **300** may include working through, filling-in and annotating content, recording sessions for archiving or later analysis, and using a visual editor **108** in 3D to create scenarios to work through, fill-in, and annotate content. Content for working through, filling in, and annotating content may include documents, models, videos, video conferencing windows, presentations, worksheets, forms, surveys, interactive 3D experiences and the like. Recording sessions for archiving or later analysis may include sharing output results of completed or annotated content.

Creating scenarios to work through, fill-in and annotate content may allow an almost unlimited range of scenarios from data rooms, to interactive 3D simulations, gaming, e-learning through to creating 3D control surfaces. Because a visual editor **108** may be built using an application system engine **102**, it also may have the ability to render content in stereo. This effectively means that any component in a scene, that may also be a 3D version of the page, may be brought into a plane view using the UI.

By toggling edit mode, a user may bring-up a 2D editor UI correctly, surrounding the zoomed and oriented 2D plane containing the component/media type. Then a user may pick and adjust all of the GUI and text elements related to the properties of an object and set-up any desired actions, as every part may also be rendered through the same pipeline. In this way, a visual editor **108** may become a real-time viewer and editor for multiple, simultaneous users of interactive VR and AR workspaces, including contention management for multi-user editing and broadcasting of edits.

In embodiments, a collaborative 3D workspace **300** may include AR/VR specific properties for positioning and tools. AR/VR specific properties for positioning and tools may include a capability to allow users to define items in a scene that may be moved, auto-placement settings, placement restrictions and how items should be grouped. For example, a user may define, in the editable properties that a series of elements are grouped and will auto-place themselves on a flat surface visible to at least 50% of participants. A user may also specify a perimeter by providing a minimum radius the center of an item can be in relation to a perimeter of another item. This may prevent items that can be moved interactively from being moved too closely together.

Referring to FIG. 2D, some objects **302** may be displayed only on flat surfaces **304**, some on vertical surfaces **306**, some in free space **308** and some as a heads-up display (HUD) or directly inside the viewport of each user, the latter being similar to how a traditional page element might be placed **310**. For example, a set of participants may look at the same table. In this example, the components in the table

may maintain their relative positions as they transition into view and keep their relative positions and scales on the table relative to wherever the participants may look.

It is noted that when a first user chooses to use a UI zoom to zoom-into a specific component, the UI zoom of a first user may, in embodiments, not move a component for other users. The object zoomed-into by a first user and their annotations may continue to appear in-place for other users. It is also noted that when a first user edits a component in UI Zoom mode, the component may stay in place for other users, however a colored halo may be displayed to other users to indicate that the editing functionality is currently locked by another user (e.g., the first user).

An advanced collaborative 3D workspace group property may have the ability to specify content density constraints for sub-groups. Specifying content density constraints for sub-groups may allow a user to specify different components to be displayed from the subgroup based on the distance of a user from an object. In embodiments, specifying content density constraints for sub-groups may be used to manage levels of detail. In an example, a large title and date of a document may be displayed when viewed at a distance by a user, then the full document may be displayed as the user zooms “closer” to the document. In embodiments, each distance level in the example may have a built-in hysteresis value. A built-in hysteresis value may be based on a ratio of distances. Basing a built-in hysteresis value on a ratio of distances may prevent different component representations from performing undesired movement, such as bouncing or popping.

As depicted in FIG. 2E, some objects **302**, such as documents and spreadsheets, may come with their own pre-set content density settings, such as large title and date, at a distance content density settings. In contrast to this example, user-defined components may be more sophisticated, such as creating a news article, which may be a title and date at a “far” distance **312**, a title and thumbnail at the “mid” distance **314**, and then the full article and images at a “close” distance **316**.

As depicted by the example in FIG. 2F, one or more avatar components may be placed into a page, such as the 3D scene depicted in FIG. 2D, and configured using AR/VR specific positioning properties. The one or more avatar components could then be set to lock onto a horizontal or vertical surface, or to float in mid-air in the example. However, the most likely setting may be as a HUD element **310**. Continuing with this example, each user may receive their own instance of an avatar at runtime. Each version of the avatar may be fully configured for the preferences, language and choice of tasks of each user. The possible choices may be defined at design time. In this example, inside a user’s personal view, a user may have their own avatar that is able to help them with tasks related to the current custom scenario. These tasks might be intended to help the user with research on a topic, configure parameters on a simulation, to perform editing or data entry tasks or any other interactive scenario which can be driven with a set of avatar natural language processing (NLP) skills, such as vocabulary, commands, and process workflow NLP skills and either local editor capabilities or cloud service lookups.

Continuing reference to the example depicted in FIG. 2F, when an object **302** is selected then the avatar **318** may provide the state of current workflow (B) and the current options for the user workflow (C) that may list of types of verbs a user can use to perform tasks. In embodiments, the list of types of verbs may not need to be a comprehensive list of verbs or verb synonyms. Similarly, the list of types of

verbs may not be the set of all possible interactive skills or workflows a user may access by voice. However, the list of types of verbs may make discovery for the current workflow possible for users who are interacting with voice commands.

If a VR or AR head mounted system supports eye tracking, then it may be possible to use the avatar voice recognition in combination with eye tracking to provide selection capabilities. For example, looking at a view may provide a cursor when the user fixes their vision, for example by staring at a point on the view. Then, if the point is located on a particular component, a user speaking “select” or another appropriate command structures, such as “please pick,” or “select the item,” may indicate a command to select the particular component. It may also be possible that the user may choose to go directly to using action verbs or skills, if the user is already familiar with the types of command verbs and skills that may be possible. The ability to go directly to using action verbs or skills may be a shortcut for expert users.

In embodiments, a command structure may include command verbs. Command verbs may be keywords. Command verbs may trigger a behavior at a given point in a skill workflow. Some command structures may include synonyms for command verbs.

In embodiments, skill workflows may be voice applications. A skill workflow may have a flexible set of stages that describe the workflow of a user. Each stage may require different values to be supplied by a user and provide different command verbs a user may use to change the tasks performed by the workflow and stage of the workflow the user is on. In embodiments, skill workflows may include rules. Rules may determine how a user may transition through workflow stages and ensure required values are specified or defaulted. FIG. 2G depicts a three-stage example of a high-level recipe workflow. In a first stage **320**, a user selects a recipe skill then moves to stage **1**. Stage **1** requires a user to input values for recipe name and quantity. In stage **1** a user is also prompted to enter allergies as an optional input value. The workflow then proceeds to stage **2** **322**. In stage **2** **322**, a user is presented a shortlist of recipes and an option to accept their name or number. In stage **2** **322** a user is also presented with a command option to go back in the workflow to change previous choices once a recipe is selected. The workflow then proceeds to stage **3** **324**. In stage **3** **324**, a user progresses through being presented with the recipe instructions, including commands such as next instruction, repeat instruction, start again and the like.

In embodiments, it may be possible to render a scene in stereo, as all visual elements in a scene tree may pass through a unified 3D pipeline. Any positional data from head mounted or phone tracking in the real world may also be supplied to the engine to modify the rendering viewport, so as to match the correct orientation, position, and field of view.

In embodiments, the ability to pick strategies may be provided to locate objects in a 3D scene using only a 2D x/y coordinate relative to a current view. This may allow users to interact with a GUI by clicking or pointing at it. Pointing may also include a direction vector which may be raycast until an item in the scene is hit.

In embodiments, a 3D UI may utilize a minimal transformation strategy to accommodate that the scene has not changed and the projection may only have moved in the X-axis. Primarily a CPU side scene tree transformation may be simplified with a reduced transformation space for each eye and as a special GUI optimization a 3D UI may keep items for each eye that exist on the zero 2D plane as they

have identical vertexes. The zero plane is important, as this is the most comfortable zone for users. As such, the zero plane may be set up at the focal distance for the given stereo hardware and provides perfect convergence. For example, tasks such as reading text and interacting with forms are ideally placed at the zero plane.

In embodiments, a 3D UI may provide a mechanism to render binocular (stereo) views from a single stream of OpenGL rendering commands. Vertex, tessellation, and geometry shaders may output two positions for each vertex corresponding to the two eye views. The positions from each view may be sent to different viewports or layers or both. FIGS. 23A and 23B depict examples of morph vertex animation performed by a GPU.

In embodiments, a 3D UI may support 3D input and user interfaces. A 3D UI may use hand proximity display menus and head proximity change density to support 3D input and user interfaces. Hand proximity display menus may respond to hand motions for example, as a hand nears an object, menus may radiate, such as when a video thumbnail may show play, rewind and stop options. Head proximity change density may respond to head motions for example, as a head gets closer to a screen the amount of content increases or as a head moves away from a screen key content enlarges, while minor content bubbles off screen and content reflows.

In embodiments, a 3D UI may support content as a virtual window. For example, as a head moves relative to a screen, the screen may act like a window. In this example, the content may be organized in concentric rectangles around the center, so a user can see more layers by moving closer or moving relative to the screen. A user may also use a vertical hand and peeking gesture to reduce the amount of head movement required.

In embodiments, a 3D UI may support hysteresis for content density. Hysteresis for content density may reduce the flickering of information between content zones using hysteresis to have overlapping close and far thresholds. This may give a user much more control while not penalizing small, natural movements of the body. Because of the parallax effects, content zones may need to have hysteresis in the horizontal and vertical axis, as well as in depth. If the space on the left of the screen is constrained, they may have metadata indicating they can reform on the right-hand side where there is more space, or even to shift other content to the side.

In embodiments, a 3D UI may support content density flexible reflow. In an example, content density flexible reflow may use a non-orthogonal camera to make the computer monitor act like a window. Then as a user approaches, distant content may appear to the human eye to stay the same size, but will occupy fewer square pixels on the monitor. Because of parallax, this content, which is distant, will move left, right and up and down from center based on the user's head position. FIG. 24 depicts examples of content density.

Content density flexible reflow is, therefore, more complex than traditional reflow and the proposed method uses information nodes that may be connected in a tree structure. These information nodes may be fitted and flowed into content zones that may be altered based on user orientation and distance from the screen. It is noted that nodes may not pop in and out of existence as the content zones will. Nodes may be animated if their position changes, or fade in and out of existence to provide the user a much clearer visual experience.

In embodiments, a tree structure for nodes may include a 'replace' tag. A 'replace' tag may allow a lower node to

replace a parent node rather than augmenting it. For example, a news heading node may contain a sub-node which is a paragraph summary. Then the paragraph summary contains two sub-nodes, a fuller description and a photo. Because it is redundant to show the summary and the fuller article, the summary node will have the 'replace' tag set.

In embodiments, a 3D UI may include ghost hands, a user interface paradigm for using 2D and 3D user interfaces with fingertips. Ghost hands may be used in both desktop and VR scenarios. For example, one or two ghost hands may be mounted in the bezel of a laptop or tablet where a 3D camera is mounted onto a stereo headset facing outward. Existing skeletal tracking solutions have problems where the final, inverse kinematic phase, which attempts to rectify the skeleton, will almost always reposition joints including the final position of the fingertips. Thus, displayed joints are almost always pushed off center and often outside of the silhouette. This may cause inaccuracy and/or jitter between frames. Hands which are only partially in view, as depicted in FIG. 25, are a very common occurrence when the hand is close to the camera and filling a lot of the view. This seriously breaks the stability of the skeletal approach making it very difficult to use for interaction. Users do not get clear feedback on why the hand tracking is failing as they are still clearly holding their hand up in front of the camera frustum (which may be a result of only a partial match). Accessibility to the bottom of the screen may, thus, be impossible with this approach.

Equally, falling back to 'hand center' approaches, as depicted in FIGS. 26A and 26B, is also broken as the external shape of the hand changes a lot, and finding a stable center of the blob for a 'cursor' is challenging—even selecting from 4 or 5 items horizontally is difficult to do with a 640×480 depth buffer. In particular, closing the fingers of the hand into a fist, as depicted in FIGS. 27A and 27B, has a number of problems for skeletal systems where the center of the hand moves, and finger bones are left in an ambiguous position. Children's hands, endomorphic versus ectomorphic hand types, and long sleeves all provide specific challenges for skeletal trackers often resulting in quite noisy results. Ghost hands may provide a machine vision fingertip tracking capability and a gesture library which can detect taps, wiggles, swipes and dwells. Ghost hands may also have a UI library for visualizing the hands and fingertips of users' hands and fingertips as if seen in a transparent mirror. Ghost hands may also show the state of fingertips. Ghost hands may allow for sub pixel accuracy, per frame accuracy with no lag, and very natural feedback for users.

In embodiments, ghost hands may be configured to overlay a transparent view of raw depth data that may have been thresholded on the view of the users' view, for example on screen or in stereo corrected view per eye. Thresholding may allow for a normal range of hand motion near the camera, but should fall off visually as users get too far away so that users can be guided back into the correct distance from the camera.

In embodiments, a cursor or crosshairs may be presented to a user on each of the detected digit tips at a smooth 60 fps. Cursors or crosshairs presented to a user in this way may be referred to as fingertip cursors. Fingertip cursors may be used for interaction with 2D and 3D user interface elements and may have icon and color indications of their state.

For stereo views, a flattened hand visualization may converge at the exact same depth as a heads-up display (HUD) so there is no ambiguity in picking the interface elements. In a full 3D UI, a z buffer may be used to suitably

distort the view for one or both of the eye views depending on how the z buffer view has been generated. For example, structured light or stereopsis based techniques will have one eye already correct, time of flight may require distorting each view by half of the ocular distance.

Because VR currently has a fixed projection/display distance, close range focus cannot be used as a cue. Close range stereo is not ideal for the human visual system to align hands with objects in a 3D world. One solution is to show the fingertip cursors converging on the UI HUD plane while the hand/fingers retain their stereo nature. The benefit of keeping the fingertip cursors and UI converged is that the usability will be much better and more natural.

In embodiments, ghost hands may be configured to support virtual reality use cases, as desktop use cases will often have hands reaching up into the view with potentially only the fingertips visible, making this approach ideal. Simplicity and repeatability are critical for user interface interactions. With ghost hands and a z buffer (640×480), it should be possible to accurately pick one of 26 items horizontally on a screen, such as keys on a piano keyboard, at a repeatability of over 98%. A clear user mental model and immediate feedback from a UI is required for discoverability and usability. Ghost hands make it easy for users to understand how they interact. The ability to have fingertips low on the view makes it easy to access the whole user interface.

In embodiments, ghost hands may be configured to enable fingertip gestures. Fingertip gestures may include hover, dwell, tap and point metaphors, swipes, dwell to lock/select and then drag, finger to drag/paint and gravity well. Swipes over an active area may be passed to the underlying object. Dwell to lock/select and then drag may include shaking or moving out of view to, for example, show a tether line while locked/selected. In a gravity-simulating gestural interface, a hover cursor may snap to the nearest UI element as if there is a gravity effect, for example, the closer to the UI the more tightly it snaps. This may give the user additional support in picking UI elements where they are required to tap or tilt the finger which may move the cursor. This effect may be disabled when in other modes like dragging or painting.

In embodiments, fingertip gestures may include five or four digit gestures. Five or four digit gestures may include open hand to rotate, panning left, right, up and down and push and pull into screen gestures. Detecting an open hand to rotate gesture may include detecting four or five digits in a high five pose. When a change in rotation is detected, the 3D UI may rotate the locked/selected item in the z axis. Panning left, right, up and down with four fingers in a row may act as x, y scrolling/rotating gesture for a locked/selected item. Push and pull into a screen may act as a modifier gesture on a locked/selected item, for example two common modifiers would be rotating around the x axis or scaling in overall size.

In embodiments, a 3D UI may include halo fingertip detection. A traditional fingertip detector using contour tracking is good at detecting fingertips, but has some computational challenges when finding the center of the fingertip. In contrast, a halo fingertip detection approach may be kernel based. In this way, the 3D UI may take advantage of a GPU, with a core motivation to generate extremely accurate and stable centroids for fingertips. Halo fingertip detection may threshold the depth data so hands are the only likely objects in the view, as depicted in FIG. 28A.

Halo fingertip detection may run a modified morphological operator to erode edges based on curvature, for example using z value differences a 3D UI may estimate a normal

vector, hence detect edges of fingers and ensure that even fingers held tightly side by side will be separated.

Halo fingertip detection may run a modified morphological operator to infill holes and average z values. Halo fingertip detection may run a coloring fill over the remaining blobs so overlapping hands can be separated and identified, as depicted in FIG. 28B.

In embodiments, halo fingertip detection may run a sparse kernel consisting of a series of concentric rings. This may be used as a fast first pass to identify the best fingertip candidates at a range of scales, as the finger will appear larger in the camera view the closer it is to the camera. These rings may be looking for an occupancy pattern where the internal rings are fully occupied and then a subsequent ring has a halo with very few values, checking only the blob color from the center rings, for example treating other blobs as non-occupied pixels, as depicted in FIG. 28C. From the list of best matches, a full kernel may be run over a fingertip to establish the best fit circle which covers the tip of the finger. The pixels inside the circle may be averaged for x, y and z values. At a midrange from the camera with a 640×480 buffer, a fingertip is roughly 60×80=4800 pixels in size. This may generate a centroid which may be accurate to a 5th of a pixel on a 640×480 buffer with strong spatial coherency between frames as the rings may keep a constant fit at the tip despite noise in the silhouette. The occupancy of an outer ring at 2× radius of the fit of the fingertip may be used to determine the angle in 2D space of the finger, i.e., the compass angle. The z sample from a boundary may be used to determine the z gradient of the fingertip, from finger to tip, to determine if it is facing into, flat, or out of the view. Some ‘fingers’ may be rejected if the radius and the z distance are outside the scale expected from human child/adult fingertips. Temporal tracking with a simple motion predictor is used to identify fingertips between frames to ensure a consistent ID number is passed back to the underlying software, as depicted in FIG. 28D. ID numbers may be recycled when a finger is not detected for at least 2 frames. This may feed into a gesture processing system.

The following is an example of Halo fingertip pseudo code—

```

Function process_fingertips_per frame( )
{
  buffer = Read_z_data( )
  z_buffer = pack_texture(buffer)
  gpu_buffer = transfer_to_GPU(z_buffer)
  threshold(gpu_buffer, maximum_z_distance)
  z_erode(gpu_buffer, z_delta_erode_ratio)
  z_fill_smooth(gpu_buffer, kernel_size, amount)
  color_blob_id(gpu_buffer, minimum_blob_pixels)
  candidate_list = sparse_halo(gpu_buffer)
  loop (candidate_list)
  {
    vector3_position_list, vector3_angle_list =
      fit_halo(candidate_list_entry)
  }
}
Function sparse_halo(gpu_buffer)
{
  foreach_pixel(gpu_buffer)
  {
    center_pixel = pixels(center_vector)
    foreach_ring( )
    {
      z_occupancy[ring] = pixels(sample_vectors)
    }
    if z_occupancy[inner_rings] == high &&
      z_occupancy-outer_rings == low
  }
}

```

-continued

```

        candidate_list += center_vector
    }
    }
    return candidate_list
}
Function fit_halo(vector candidate_list_entry)
{
    foreach_pixel_in_search_space(gpu_buffer)
    {
        foreach_ring( )
        {
            z_occupancy[ring] = pixels(full_circle)
        }
        if z_occupancy[inner_rings] == high &&
            z_occupancy[outer_rings] == low &&
            ring_scale > min && ring_scale < max
        {
            vector3_position_list += precise_center_vector
            vector3_angle_list += angle_vector
        }
    }
    return vector3_position_list, vector3_angle_list
}

```

In an example of halo fingertip detection and referring to FIG. 29A, fingers tilted forward (tilt) or currently being moved toward the screen in the Z-axis (tap) will show with a highlight 2902 indicating they are activating the piano key they are directly above. In another example of halo finger tracking and as depicted in FIG. 29B, by tracking fingertips the same camera used to scan the world (for augmented reality scenarios) can be used. A user does not need to have their hand completely in view. With the ghost hands approach, fingertips may show up clearly, while traditional skeletal tracking faces ambiguous skeletal poses of the back of the hand.

In embodiments, the application system 100 may include the engine 102. The engine 102 may connect to an editor and engine editor and runtime infrastructure 104, the declarative language 140, cloud services 142, third party modules 146, and the visual editor 108. The editor and engine editor and runtime infrastructure 104 may connect to the engine user interface (UI) 106. The engine UI 106 may connect to a viewer/portal 144 and an avatar interaction engine 148.

As depicted throughout this disclosure and in some embodiments, the application system 100 may include the engine 102 that unifies the creation, editing and deployment of an application across endpoint devices, including endpoint devices that run heterogeneous operating systems. Thus, an app created in the application system 100 can automatically operate, without creation of separate versions, on different mobile operating platforms, such as Android™ and IOS™ platforms.

In embodiments, the engine 102 may be configured to support a multi-user infrastructure by which, for example, different users of the engine may edit a scene tree description 124 for an application 150 or otherwise collaborate to create an application. Each user may edit the scene tree description 124 for the application simultaneously with other users of the visual editor 108. In embodiments, a user may edit the scene tree description 124 for the application 150 simultaneously with other users of the visual editor 108 or users of the runtime of the application 150. Any rendered depictions (e.g., simulations) of the behavior of objects in the application 150 may appear the same to all users of the visual editor 108 and/or of the runtime infrastructure 104 of the application 150.

In embodiments, the engine 102 may share the editor and runtime infrastructure 104 for the code that implements the

application 150. The same engine 102 may be used for both editing and running the application, using the same editor and runtime infrastructure 104.

In embodiments, the engine 102 may include a visual code editing environment, also referred to as a visual editor 108 throughout this disclosure. The visual editor 108 may allow a developer to code high-level application functions, including how an application 150 will use the CPU/GPU of an endpoint device that runs the application 150, to enable optimization of the performance of the application 150.

In embodiments, the engine 102 may include a gaming engine 119 to handle machine code across different operating system platforms, within the same editor interface. The engine 102 may include a plug-in system 121, such as a JavaScript plug-in system, a visual editor 108, a script layer 134 and additional engines, such as a serialization engine, browser engine 226 or 3D engine 164, to simulate and run of code developed using the application system 100. The engine 102 may include a shared simulation of the runtime behavior of the application 150 that is being edited by the visual editor 108.

In embodiments, the engine 102 may include a declarative language (also referred to as a dynamic language 140 in this disclosure). In embodiments, the declarative language 140 may describe a scene tree. The declarative language may describe a scene tree using a scene tree description 124. The scene tree description 124 may specify the page layout of an application 150 and the structure of interactions among application elements in response to user input. The engine 102 may include a coding environment, such as a content and application creator 109, that includes the dynamic language 140, in which the runtime and the editor for the application may be compiled by an LLVM compiler 136. The engine 102 may include the ability to express logic for application behavior, as well as presentation layer layouts of visual elements for the application 150, in the same declarative language 140. The declarative language 140 may include object classes and methods. Object classes and methods may allow a developer to specify conversational interfaces and natural language endpoints. Conversational interfaces and natural language endpoints may be used as inputs to create an emotionally responsive avatar for integration into a system that uses the avatar. The conversational interface and natural language endpoints may be received as inputs and the avatar may be created using an avatar interaction and rendering engine 148 of the engine 102.

In embodiments, the engine 102 may use or include a domain-specific language. The same domain-specific language may be used for the editor and runtime infrastructure 104 and file management for applications 150 developed using the application system 100.

In embodiments, the engine 102 may include an application development language. The application development language may include named states and support an unlimited number of named states. Named states may be added to an object or object class which may encapsulate properties or methods. The application development language may be extended with new object classes, methods, properties and events. Methods, properties and events may expose device features across devices. Devices may use the same or different operating system platforms. The application development language may include a domain-specific language for visual experience creation.

In embodiments, the engine 102 may include a development environment, such as a content and application creator 109. The content and application creator 109 may include or connect to a viewer/portal 144. A viewer/portal may be a

private portal. The private portal may publish applications **150** from the development environment without requiring deployment through an app store, for example. The content and application creator **109** may connect to the engine **102** through an engine API **114**.

In embodiments, the engine **102** may include an avatar class operating with JavaScript, such as within an avatar interaction and rendering engine **148**. The avatar interaction and rendering engine **148** may specify parameters for speech synthesis and 3D animation for an avatar.

In embodiments, the engine **102** may include or connect to an engine user interface **106**. The engine user interface **106** may enable non-technical users to specify variations **190**. Variations **190** may be variations of objects visually presented in applications **150**.

In embodiments, the engine **102** user interface **106** may allow users to manage the state of one or more objects that may be used to create a visual presentation for an application **150**. The engine user interface **106** may also include one or more interfaces for handling input parameters for 3D content and other input parameters. Other input parameters may include content density parameters, hand proximity display parameters, head proximity change density parameters, content as a virtual window parameter, hysteresis for content density parameters, content density parameters and the like. In embodiments, the engine user interface **106** may include support for 3D content generation. 3D content generation may be generated by a 3D engine **164**. The engine user interface **106** may include support for 3D content generation may include the ability for a user of the engine user interface **106** to hot key, in order to identify the object and properties of a 3D object for the application. The 3D engine **164** may include an editor for handling 3D machine vision input. 3D machine vision input may manage color information and information relating to a distance from a defined point in space. The engine user interface **106** may also include an application simulation user interface. The application simulation user interface may share the infrastructure and engine for the code that implements an application **150**. The application simulation interface may connect to the visual editor **108**.

In embodiments, the editor and runtime infrastructure **104** and the engine **102** may allow for the creation, editing and running of an application **150** that includes both 2D and 3D elements. The editor and runtime infrastructure **104** may use a hybrid scene tree description **124** that includes 2D and 3D elements that may be rendered, composited and interacted within the same visual scene of an application **150**.

The editor and runtime infrastructure **104** may allow for a scene tree description for an application **150** to be edited simultaneously by multiple users of the editor and runtime infrastructure **104** of the application **150**, such that rendered simulations may appear the same to all users.

In embodiments, the editor and runtime infrastructure **104** may connect to the visual editor **108**. The visual editor **108** may allow a developer to code high-level application functions and may define how an application **150** may use the CPU/GPU of an endpoint device that runs the application **150**, to enable optimization of application performance.

In embodiments, the engine user interface **106** may include support for the simulation of an application **150**. The engine user interface **106** may share the editor and runtime infrastructure **104** and engine **102** for the code that implements the application **150**. The editor and runtime infrastructure **104** may include a visual editor **108** that may use the same engine **102** for editing and running the application **150**.

In embodiments, the engine **102** may include a gaming engine **119**. The gaming engine **119** may handle machine code across different operating system platforms within the same engine user interface **106**. The editor and runtime infrastructure **104** may include a plug-in system **121**, for example, a JavaScript plug-in system, a visual editor **108**, a script layer **134** and an engine, such as a serialization engine **112**, for simulation and running of code developed using the application system **100**.

In embodiments, the visual editor **108** may include a shared editing environment. The shared editing environment may enable real time, multi-user, simultaneous development, including shared simulation of the runtime behavior of the application **150** that is being edited. The shared editing environment may be synchronized by a multi-user layer sync application and asset system **120**. The visual editor **108** may include support for the dynamic language **140**, private portal, editing engine, object classes, 3D content, 3D content generation user interface and hybrid 2D and 3D scene trees as described previously in this disclosure.

In embodiments, an application system **100** is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems. In embodiments, an application system **100** is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users. In embodiments, an application system **100** is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application. In embodiments, an application system **100** is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a visual code editing environment that uses the same engine for editing and running an application. In embodiments, an application system **100** is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a visual code editing environment wherein a developer can code high-level application functions and can code how an application will use the CPU/GPU of an endpoint device that runs the application to enable optimization of application performance. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a visual code editing environment that uses a gaming engine to handle machine code across different operating system platforms within the same editor interface. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a JavaScript Plug-In system, an editor, a script layer and an engine for simulation and running of code developed using the system. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a

shared editing environment enabling real time, multi-user, simultaneous development, including shared simulation of the runtime behavior of an application that is being edited. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a declarative language that is used to describe a scene tree that specifies the page layout of an application and the structure of interactions among application elements in response to user input. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a coding environment with a declarative language in which the runtime and the editor for an application are compiled by LLVM. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having the ability to express logic for application behavior and presentation layer layouts of visual elements for the application in the same declarative language. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having the same domain-specific language for an editor, a runtime and file management for applications developed using the system. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a development language with unlimited named states that can be added to an object or object class which can encapsulate properties or methods. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having an application development language designed to be extended with new object classes with methods, properties and events which can expose device features across devices using different operating system platforms. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a domain-specific language for visual experience creation. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a private portal within a development environment to publish applications from the development environment without requiring deployment through an app store. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having an editing engine within a development environment that includes an avatar class operating with JavaScript, wherein the avatar class specifies parameter for speech synthesis and 3D animation for an avatar. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a declarative language with object classes and methods that allow a developer to specify conversational interfaces and natural language endpoints to create an emotionally responsive avatar for integration into

a system that uses the avatar. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a user interface that enables non-technical users to specify variations of visually presented objects in an application. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a variations user interface layer that allows users to manage the state of one or more objects that are used to create a visual presentation for an application. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having one or more interfaces for handling input parameters for 3D content, the input parameters selected from the group consisting of content density, hand proximity display, head proximity change density, content as a virtual window, hysteresis for content density, and content density. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having a user interface for 3D content generation including the ability to hot key to identify the object and properties of a 3D object for an application. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having an engine and editor for handling 3D machine vision input that manages color information and information relating to distance from a defined point in space. In embodiments, an application system is provided having an engine that unifies the creation, editing and deployment of an application across endpoint devices that run heterogeneous operating systems and having an editor and engine for creating, editing and running an application that has 2D and 3D elements, wherein the editor uses a hybrid scene tree system that includes 2D and 3D elements that can be rendered, composited and interacted within the same visual scene.

In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a visual code editing environment that uses the same engine for editing and running an application. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a visual code editing environment wherein a

developer can code high-level application functions and can code how an application will use the CPU/GPU of an endpoint device that runs the application to enable optimization of application performance. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a visual code editing environment that uses a gaming engine to handle machine code across different operating system platforms within the same editor interface. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a JavaScript Plug-In system, an editor, a script layer and an engine for simulation and running of code developed using the system. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a shared editing environment enabling real time, multi-user, simultaneous development, including shared simulation of the runtime behavior of an application that is being edited. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a declarative language that is used to describe a scene tree that specifies the page layout of an application and the structure of interactions among application elements in response to user input. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a coding environment with a declarative language in which the runtime and the editor for an application are compiled by LLVM. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having the ability to express logic for application behavior and presentation layer layouts of visual elements for the application in the same declarative language. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having the same domain-specific language for an editor, a runtime and file management for applications developed using the system. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a development language with unlimited named states that can be added to an object or object class which can encapsulate properties or methods. In embodiments, an application system is provided

having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having an application development language designed to be extended with new object classes with methods, properties and events which can expose device features across devices using different operating system platforms. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a domain-specific language for visual experience creation. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a private portal within a development environment to publish applications from the development environment without requiring deployment through an app store. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having an editing engine within a development environment that includes an avatar class operating with JavaScript, wherein the avatar class specifies parameter for speech synthesis and 3D animation for an avatar. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a declarative language with object classes and methods that allow a developer to specify conversational interfaces and natural language endpoints to create an emotionally responsive avatar for integration into a system that uses the avatar. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a user interface that enables non-technical users to specify variations of visually presented objects in an application. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a variations user interface layer that allows users to manage the state of one or more objects that are used to create a visual presentation for an application. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having one or more interfaces for handling input parameters for 3D content, the input parameters selected from the group consisting of content density, hand proximity display, head proximity change density, content as a virtual window, hysteresis for content density, and content density. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an

application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having a user interface for 3D content generation including the ability to hot key to identify the object and properties of a 3D object for an application. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having an engine and editor for handling 3D machine vision input that manages color information and information relating to distance from a defined point in space. In embodiments, an application system is provided having a multi-user infrastructure that allows the editor to edit a scene tree for an application simultaneously with other users of the editor or users of the runtime of the application such that rendered simulations appear the same to all users and having an editor and engine for creating, editing and running an application that has 2D and 3D elements, wherein the editor uses a hybrid scene tree system that includes 2D and 3D elements that can be rendered, composited and interacted within the same visual scene.

In embodiments, an application system is provided having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application. In embodiments, an application system is provided having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application and having a visual code editing environment that uses the same engine for editing and running an application. In embodiments, an application system is provided having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application and having a visual code editing environment wherein a developer can code high-level application functions and can code how an application will use the CPU/GPU of an endpoint device that runs the application to enable optimization of application performance. In embodiments, an application system is provided having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application and having a visual code editing environment that uses a gaming engine to handle machine code across different operating system platforms within the same editor interface. In embodiments, an application system is provided having a user interface for simulation of an application that shares the infrastructure and engine for the code that implements the application and having a JavaScript Plug-In system, an editor, a script layer and an engine for simulation and running of code developed using the system.

In embodiments, creating, sharing, and managing digital content, such as for experiencing on a plurality of different digital device end points in a network may be accomplished by a system that incorporates a visual editing environment with a code execution environment that work together to enable at least one of creation, delivery, and editing of a digital asset during runtime of the asset. The combination may further enable a plurality of end users using different devices to concurrently experience the same behavior of the digital asset during its creation and its editing. In embodiments, the visual editing environment may enable a developer to create and edit code controlling a digital content asset. In embodiments, the code execution engine may operate in the visual editing environment, such as on the created code to control execution of hardware elements,

such as a hardware infrastructure element that enables utilization of the digital content asset.

In embodiments, a system that combines the visual editing environment and the code execution engine to enable creation-time and editing-time consistent behavior across different devices may include the visual editing environment interacting with the code execution engine during, for example creation of digital content asset management code. In embodiments, a user interacting with the visual editing environment may, such as through graphical manipulation of a digital content asset, may effectively be producing code that the code execution engine responds to by executing it to cause visual effects in the visual editing environment, and the like. The code execution engine may be configured to respond to the produced code with data and graphic manipulation functions for each different device based, for example on device identification information, such as the type of operating system that the different devices are executing, and the like.

In embodiments, the code execution engine that works cooperatively with the visual editing environment to facilitate presenting consistent behavior of the digital content asset on different devices may be the same code execution engine that supports runtime operation of an application that uses the digital content asset control code generated with the visual editing environment. In embodiments, the code execution engine may operate in association with an executable container that may include digital content assets, such as to enable viewing the digital content assets with a viewer that, based at least in part on the code execution engine facilitates the consistent digital content asset behavior across the different devices, such as end point devices in a network and the like. In embodiments, the visual editing environment may facilitate multiple users and/or groups to simultaneously create code for controlling a digital content asset. The visual editing environment may benefit from the code execution engine managing utilization of a CPU, a GPU, or the like, for digital content asset and optionally general software code development. In embodiments, the visual editing environment may include capabilities for use of a gaming engine, including for coding gaming behavior of hardware and software components, including operating system and hardware platforms. In embodiments, the visual editing environment may enable creating and editing a digital content asset control application in declarative language that may facilitate specifying control of machine behavior of a device, the abstraction of input types across operating system types, the capability for control of visual presentation layer behavior of objects across a plurality of operating system platform types, and the like.

In embodiments, the visual editing environment may incorporate and/or be accessible through a user interface that may provide access to the functions of the visual editing environment including without limitation digital content asset functions of creating, editing, sharing, managing, publishing, and the like. The visual editing environment, such as through the user interface may facilitate presentation of the digital content asset and its behavior by interacting with the code execution engine, such as through code execution engine Application Programming Interface (API). In embodiments, the visual editing environment may facilitate the users of the different devices to simultaneously experience the same behavior of the digital content asset through a multi-user synchronization system that may operate as part of the visual editing environment to allow, among other things simultaneous experience, editing, and the like. In embodiments, multiple instances of the visual editing envi-

ronment may be active on a portion of the different devices and may be synchronized via the multi-user synchronization system.

In embodiments, the visual editing environment may utilize the code execution engine as fully compiled code, which may facilitate achieving the simultaneous experience of the same behavior for different devices, such as tablets and the like that may not support runtime compilation, and the like.

In embodiments, a code execution engine of a system for creating, sharing and managing digital content may control utilization of hardware resources of the different devices, such as CPUs, GPUs and the like. Utilization of, for example, CPUs of some of the different devices, such as hardware endpoint devices and the like may be controlled to facilitate the simultaneous experience of the same behavior. In an example of CPU utilization, a code execution engine may utilize a graphic drawing capability of a CPU on the devices so that the behavior of the digital asset is experienced the same on the different devices. By controlling CPU utilization as in this example, differences that may be experienced when using a CPU on a first device and a GPU on a second device to perform a graphic display operation may be avoided. In embodiments, CPU and GPU utilization control may further facilitate simultaneous experience of users on different devices by, for example, allowing for rapid deployment of digital content asset behavior code across the devices without having to customize the deployment to utilize a CPU on a first device that does not have a GPU and an available GPU on a second device.

In embodiments, the code execution engine may operate with the visual editing environment during creation, editing, and the like as well as during runtime of digital content asset code generated by use of the visual editing environment. With the same code execution engine operating during visual editing and during runtime, visual editing may result in generating code, such as digital content asset control code that can control utilization of a CPU and/or GPU, such as by generating code execution control statements. Code execution control statements may include hardware resource utilization statements that may directly control utilization of different device hardware resources, such as a CPU, GPU, and the like. In embodiments, a language used in the visual editing environment, such as a declarative language that may be described herein, may include hardware resource utilization statements that the code execution engine may execute or that may affect how the code execution engine executes code, such as executing a graphic function with a CPU even when a GPU is available on the device, and the like.

In embodiments, the code execution engine of the system for creating, sharing and managing digital content may further control utilization of hardware resources for different aspects of hardware performance, including thermal performance, battery management, and the like. In embodiments, the system may have access to instruction-level execution power and thermal performance information for different devices. Device-specific instances of the code execution engine, for example, may be represented at the instruction-level so that the impact on at least one of thermal and power performance may be determined for each instruction that may be executed by the code execution engine on the devices. The digital content asset control code created and/or edited, such as by a developer with the visual editing environment that the code execution engine will perform can be analyzed based on the power and/or thermal impact of each corresponding device-specific instruction. The result of

this analysis may be a measure of the thermal and/or power (e.g., battery demand) impact on the device so that the impact may be controlled. In embodiments, the analysis of the digital content asset control code that the code execution engine may execute may suggest specific code and/or execution control of that code, such as a specific sequence of instructions, a rate of execution of instructions, or the like that may reduce or optimize thermal performance of the device. In embodiments, optimizing thermal performance for a hardware resource of one or more different devices for which utilization may be controlled, such as a CPU, a GPU and the like may be based on computation of a thermal impact of executing a digital content asset control code set by a code execution engine. This thermal impact computation may include CPU utilization (e.g., execution rate and the like), GPU utilization, memory utilization, and the like and may be determined by the thermal impact of instructions, such as CPU instructions from the digital content asset control code, generated by a compiler. In embodiments, the thermal impact computation may include compiled instructions generated from a code execution engine performing the digital content asset control code on the device. In embodiments, thermal optimization may include minimizing temperature rise of hardware resources of a device, such as the CPU for a portion of the digital content asset control code generated by, for example, a developer in the visual editing environment, and the like. In embodiments, thermal optimization may include achieving an average temperature rise during execution of a portion of the digital content asset control code by the code execution engine. This may include allowing portions of the digital content asset control code being executed by the code execution engine to result in a temperature that exceeds an optimized average temperature, while ensuring that an average temperature rise while executing a specific portion of the digital content asset control code does not exceed the average. In embodiments, thermal optimization may include limiting a temperature rise of one or more hardware resources of a device, such as CPU, GPU, and the like from exceeding an optimized maximum temperature. This may include reducing a frequency of execution by the code execution engine where a higher frequency of execution would result in a temperature increase beyond the optimized maximum temperature.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the common behavior, over a plurality of operating systems. In embodiments, the code execution engine may govern execution of the digital content asset control code to provide the consistent user experience. A consistent user experience may include a look and feel of a user interface, a speed of control of a digital content asset, interaction elements in the user interface, and the like. A code execution engine may perform a runtime check of the operating system of a device on which the code execution engine is executing the digital content asset control code and adjust a sequence of instructions, a selection of instructions, and the like based on a result of the runtime operating system check. A code execution engine may limit selection of instructions to be generated by a compiler for each of a plurality of different operating systems so that instructions being executed on different operating systems result in a consistent user experience. In embodiments, a code execution engine may take an abstraction of a visual element from a digital content asset control code and convert it into a set of instructions that ensure a consistent user experience across different operat-

ing systems. This may take the form of graphic primitives, and the like to generate a consistent visual element for each operating system. In embodiments, an operating system native icon that is used to activate an operating system function, such as rotating a digital content asset in a graphical user interface may appear differently on different operating systems. One operating system may show a curved arrow along an edge of the digital content asset; yet another operating system may show a circular icon along a central axis of the digital content asset. The code execution engine may, instead of executing instructions that generate the native icon, may execute instructions that generate a digital content asset rotation icon that is consistent across different operating systems, independent of the native icon for this purpose that each operating system uses.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may ensure simultaneous user experience of the same behavior of the digital content asset by, for example, decoupling the language from target system-specific resources, such as a device and/or operating system specific rendering engine and the like. In embodiments, rendering actions, as may be represented by a rendering layer of functionality, may be encoded into a language used in the visual editing environment to generate digital content asset control code and the like. In this way, a digital content asset may perform the same behavior on different devices, different operating systems, and combinations thereof. In embodiments, the code execution engine may function similarly to a gaming engine in that the combination of visual editing environment language and code output therefrom with the code execution engine define behaviors at a rendering level of the digital content asset, such as 3D movements, and the like. In embodiments, this combination facilitates coding the users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior at the time that the digital content asset control code is generated in the visual editing environment. In this way, the code execution engine, which may function similarly to a gaming engine for this specific aspect, may do the underlying work of making the behavior consistent, without a developer having to consider how any target device and/or operating system may need to be controlled to generate the desired behavior.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may govern execution of the code for a consistent user experience, such as is described herein, across a plurality of mobile operating systems, including, without limitation operating systems such as IOS™, ANDROID™, WINDOWS™, and the like. Mobile operating systems may include their own look and feel, including how fundamental user interactions are performed. Governing execution of code across mobile operating systems may include adapting execution of digital content asset control code and the like so that, while the user experience may not be the same across mobile operating systems, a user of an IOS™ based device may experience the digital content asset with a look and feel that a user experiences when using other mobile applications on the device. In this way, a consistent user experience may be tailored to each individual mobile operating system so that the digital content asset may effectively appear to have the same behavior, while the underlying user interface and mobile operating system native controls may be preserved for each mobile operating system. This may involve, for example, distinguishing digital content asset control code that should be executed consistently on each mobile operating system from code that should be directed to a mobile

operating system specific user experience. In embodiments, rotational behavior of a digital content asset may be consistent across mobile operating systems, whereas the controls for rotation may be operating system specific.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may govern execution of the code for a consistent user experience, such as is described herein, across a plurality of computer operating systems, including, without limitation operating systems such as MAC™, LINUX™, WINDOWS™, and the like. Computer operating systems may include their own look and feel, including how fundamental user interactions are performed. Governing execution of code across computer operating systems may include adapting execution of digital content asset control code and the like so that, while the user experience may not be the same across computer operating systems, a user of an IOS™ based device may experience the digital content asset with a look and feel that a user experiences when using other computer applications on the device. In this way, a consistent user experience may be tailored to each individual computer operating system so that the digital content asset may effectively appear to have the same behavior, while the underlying user interface and computer operating system native controls may be preserved for each computer operating system. This may involve, for example, distinguishing digital content asset control code that should be executed consistently on each computer operating system from code that should be directed to a computer operating system specific user experience. In embodiments, rotational behavior of a digital content asset may be consistent across computer operating systems, whereas the controls for rotation may be operating system specific.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may govern execution of the code for a consistent user experience, such as is described herein, in deployments that include combinations of mobile operating systems (e.g., an IPHONE™) and a computer (e.g., a WINDOWS™ LAPTOP). In embodiments, a combination of the code execution engine and the visual editing environment may facilitate this consistent user experience across mobile, computer, and other operating systems, such as by enabling creation, delivery and editing of the digital content asset during runtime (e.g., when the code execution engine is executing digital content asset control code, and the like). In embodiments, the code execution engine may be executing code, effectively in a runtime mode, during visual editing environment operations, such as creation, delivery, editing and the like. Operation of the code execution engine on each operating system type during, for example editing of a digital content asset, may be execution of a compiled combination of the digital content asset control code that is generated during the editing action and a portion of the code execution engine that executes the generated digital content asset control code. This may contrast with generating a set of graphic manipulation commands that are delivered, optionally in real-time, from a device on which a user is editing a digital content asset, to a corresponding digital content asset viewer executing on the different operating systems. In embodiments, LLVM compilation may facilitate generation of operating system-specific sets of compiled instructions that may include a portion of the digital content asset control code and a portion of the code execution engine to perform one or more digital content asset control actions

consistently across a plurality of operating systems, including mobile operating systems, computer operating systems, and the like.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may enable control of network layer interactions for the digital content asset. In embodiments, the code execution engine may be structured with a layer that may facilitate controlling network layer interactions, such as a network layer **192** of the code execution engine. In embodiments, the code execution engine may gain control of network layer interactions for the digital content asset via network layer interaction control statements that may be output from editing and digital content asset control code generation actions within the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment network layer interaction statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, network layer interactions and the like may be controlled. In embodiments, a network layer interaction that may be controlled may include a mirroring function of the network so that network actions on a portion of the digital content asset or the like may result in mirroring of a result of a network action, such as having a portion of the digital content asset or the like being mirrored to one or more different devices that may be connected to the network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices. In embodiments, network layer interaction and other network functions that may be controlled by the code execution engine, such as is described herein and the like, may include control of wireless device interactions. In embodiments, wireless device interactions may include wireless network layer interactions. In embodiments, wireless device interactions that may be controlled by the code execution engine and the like may include Bluetooth interactions, wireless access point interactions, wireless beacon interactions, near-field communication interactions, and any other wireless device interaction, and the like.

In embodiments, similarly to control of network layer interactions, the code execution engine may, such as through execution of statements that may be output from developer activity in the visual editing environment, control communication protocol interactions for the digital content asset. In embodiments, the code execution engine may be structured with a layer that may facilitate controlling communication protocol interactions, such as a communication protocol layer of the code execution engine.

In embodiments, the code execution engine may gain control of communication protocol interactions for the digital content asset via communication protocol interaction control statements that may be output from editing and digital content asset control code generation actions within the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment communication protocol interaction statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, communication protocol interactions and the like may be controlled. In embodiments, communication protocol interactions that may be controlled may include secure protocol interactions, secure socket interactions, HTTPS interactions, serial protocol interactions, and the like. Communication protocol interactions that may be controlled may facilitate communicating with one or more different devices that may be connected to the network and may, optionally be

rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the code execution engine may gain control of browser interactions for the digital content asset via browser interaction control statements that may be output from editing and digital content asset control code generation actions within the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment browser interaction statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, browser interactions and the like may be controlled. In embodiments, browser interactions that may be controlled may include Comet interactions, HTTP streaming interactions, Ajax push interactions, reverse Ajax interactions, secure socket interactions, and HTTP server push interactions, and the like. Browser interactions that may be controlled may facilitate browser interactions of one or more different devices that may be connected via a network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the code execution engine may gain control of networking middleware for the digital content asset via networking middleware control statements that may be output from editing and digital content asset control code generation actions within the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment networking middleware statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, networking middleware and the like may be controlled. In embodiments, networking middleware that may be controlled may facilitate network interaction and the like of Raknet middleware, a gaming engine, a transport layer interaction, a UDP interaction, a TCP interaction, a 3D rendering engine, a gestural engine, a physics engine, a sound engine, an animation engine, and the like. Networking middleware that may be controlled may facilitate network interactions and the like of one or more different devices that may be connected via a network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the system for creating, sharing and managing digital content may further include functions that facilitate orchestrating components, events, response to triggers and the like for a digital content asset. In embodiments, orchestrating may include without limitation automated arrangement, coordination, and management of components, events and the like. Such orchestrating may facilitate control of aspects of the system to accomplish goals of the system, such as control of hardware infrastructure elements of different devices, delivery and editing of a digital content asset during runtime, simultaneous user experience for a plurality of users of different devices, such users experiencing the same behavior of the digital content asset and the like. In embodiments, orchestrating functionality may be enabled by a plug-in capability of the system, where an orchestrating capability may be plugged-in to the system. In embodiments, the plug-in capability of the system may be a JAVASCRIPT™ compatible plug-in system.

In embodiments, a system for creating, sharing and managing digital content is depicted in FIG. **30**. The system **3000** may comprise a visual editing environment **3002** that may enable a developer **3018** and the like creating, editing, and delivery of digital content asset control code **3004**. The

system **3000** may further comprise a code execution engine **3008** may facilitate hardware resource utilization and thermal optimization for CPUs, GPU, and the like of different devices, such as end user devices **3012** on which a digital content asset **3010** operates. The code execution engine **3008** may also control interactions, such as network layer interactions, communication protocol interactions, browser interactions, network middleware and the like. The system **3000** may further comprise an orchestration facility **3014** for orchestrating components, events, and the like associated with the digital content asset **3010**. The orchestration facility **3014** may further comprise a plug-in system **3020**. The system **3000** may further facilitate consistent user experience across a plurality of different operating systems **3022**, including mobile, computer and other operating systems.

In embodiments, a system for creating, sharing and managing digital content may include a visual editing environment, a code execution engine, a declarative language, and use of a Low-Level Virtual Machine (LLVM) compiler. The visual editing environment may enable a developer to create and edit code controlling a digital content asset in the declarative language. In embodiments, the visual editing environment may enable a developer to create and edit declarative language code controlling a digital content asset. In embodiments, the code execution engine may operate in the visual editing environment, such as on the created declarative language code to control execution of hardware elements, such as a hardware infrastructure element that enables utilization of the digital content asset.

In embodiments, a system that combines the visual editing environment and the code execution engine to enable at least creating and editing declarative language code for controlling a digital content asset may include the visual editing environment interacting with the code execution engine during, for example creation of digital content asset management code. In embodiments, a user, such as a developer interacting with the visual editing environment may, such as through graphical manipulation of a digital content asset and the like, may effectively be producing declarative language code that the code execution engine responds to by executing it to cause visual effects in the visual editing environment, and the like. The code execution engine may be configured to respond to the produced declarative language code with data and graphic manipulation functions as may be expressed in the declarative language. In embodiments, the declarative language and the code execution engine used during editing of the digital content asset in the visual editing environment are also used at runtime. In embodiments, the visual editing environment and the code execution engine are constructed from this declarative language. In embodiments, the visual editing environment, runtime code derived from the declarative language, and optionally the code execution engine are compiled for distribution using a LLVM-based compiler architecture. In embodiments, the visual editing environment makes the declarative language available to a developer to generate the digital content asset management code.

In embodiments, the code execution engine may operate with the visual editing environment during creation, editing, and the like as well as during runtime of digital content code generated by use of the visual editing environment. With the same declarative language and code execution engine operating during visual editing and during runtime, visual editing may result in generating code, such as digital content asset control code that can control execution of hardware infrastructure elements, such as by generating code execution control statements. Code execution control

statements may include hardware execution statements that may directly control hardware resources, such as a CPU, GPU, and the like that may participate in utilization of the digital content asset. In embodiments, a language used in the visual editing environment, such as a declarative language that may be described herein, may include hardware execution statements that, when compiled with an LLVM compiler, for example may be executed by the code execution engine. In embodiments, the visual editing environment output digital content asset control code may be compiled with the code execution engine, such as with an LLVM compiler-based architecture to produce runtime code. In embodiments, the compiled runtime code may be distributed as it is compiled to facilitate efficient distribution and timely execution on computing devices. In embodiments, the visual editing environment, and runtime code produced from editing, creation and other actions of a user of the visual editing environment may be bound to the code execution engine, such as through compilation, such as with an LLVM-based compiler architecture and the like. In embodiments, the code execution engine may comprise a C++ engine. The code execution engine may perform execution of C++ code; the code execution engine may be coded in C++; the code execution engine may facilitate execution of C++ code as well as other code types, such as the declarative language type described herein, and the like.

In embodiments, the system for creating, sharing and managing digital content may be adapted so that a digital content asset may be edited and the like via a visual editing environment of the system, compiled, linked to the code execution engine and executed, such as by the code execution engine without depending on comparable tools (e.g., editors, compilers, execution engines, and other code generation and execution tools), and the like. The digital content asset may be compiled, linked to the code execution engine and executed by the code execution engine independent of support by tools outside of the visual editing environment and the code execution engine. In embodiments, the system may cover the domain of tools required to build and deploy digital content, such as a digital content asset and the like, while maintaining integration of elements and simplicity of configuration, use, portability, update, extension, and the like. In embodiments, the code execution engine may auto compile, effectively compiling itself. In embodiments, the engine and declarative language may be automatically compiled at runtime. The declarative language for controlling a digital content asset may be instantiated upon loading it into an execution environment. The code execution engine may be a fully compiled binary that is compiled at runtime.

In embodiments, a system that uses declarative code to control a digital content asset and a code execution engine to enable creation-time and editing-time consistent behavior across different devices may include the visual editing environment interacting with the code execution engine during, for example creation of digital content asset management code, as described herein with declarative language and the like. In embodiments, a user interacting with the visual editing environment, such as through graphical manipulation of a digital content asset, may effectively be producing code that the code execution engine responds to by executing it to cause visual effects in the visual editing environment, the runtime environment, and the like. The code execution engine may be configured to respond to the declarative language code with graphic manipulation functions for each different device based, for example on device identification information, such as the type of operating system that the different devices are executing, and the like.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content through use of a declarative language and the same language and code execution engine being used during editing and runtime may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the same behavior, over a plurality of devices. In embodiments, the code execution engine may govern how the declarative language is executed to provide the consistent user experience. A consistent user experience may include a look and feel of a user interface, a speed of control of a digital content asset, interaction elements in the user interface, and the like. In embodiments, a code execution engine may take a declarative language statement for a digital content asset and convert it into a set of instructions that ensure a consistent user experience across different operating systems. This may take the form of graphic primitives, and the like to generate a consistent visual element for each operating system.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may ensure simultaneous user experience of the same behavior of the digital content asset by, for example, decoupling the declarative language from target system-specific resources, such as a device and/or operating system specific rendering engine and the like. In embodiments, rendering actions, as may be represented by a rendering layer of functionality, may be encoded into the declarative language used in the visual editing environment to generate digital content asset control code and the like. In this way, a digital content asset may perform the same behavior on different devices, different operating systems, and combinations thereof. In embodiments, the code execution engine may function similarly to a gaming engine in that the combination of the declarative language, and optionally runtime code output therefrom with the code execution engine define behaviors at a rendering level of the digital content asset, such as 3D movements, and the like. In embodiments, this combination facilitates coding the users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior with the declarative language at the time that the digital content asset is created, edited and the like, such as in the visual editing environment. In this way, the code execution engine, which may function similarly to a gaming engine for this specific aspect, may do the underlying work of making the behavior consistent, without a developer having to consider how any target device and/or operating system may need to be controlled to generate the desired behavior.

In embodiments, a code execution engine of a system for creating, sharing and managing digital content through use of a declarative language may control utilization of hardware resources of a plurality of different devices, such as CPUs, GPUs and the like. Utilization of, for example, CPUs of some of the different devices, such as hardware endpoint devices and the like may be controlled to facilitate users of different devices experiencing the same behavior. In embodiments, the code execution engine may operate with the visual editing environment during creation, editing, and the like as well as during runtime of digital content asset code generated by use of the visual editing environment. With the same code execution engine operating during visual editing and during runtime, and the same declarative language being utilized by the visual editing environment and the code execution engine, visual editing may result in generating code, such as digital content asset control code that can control utilization of a CPU and/or GPU, such as by

generating code execution control statements. Code execution control statements may include hardware resource utilization statements that may directly control utilization of different device hardware resources, such as a CPU, GPU, and the like. In embodiments, a language used in the visual editing environment, such as a declarative language, may include hardware resource utilization statements that the code execution engine may execute or that may affect how the code execution engine executes code, such as executing a graphic function with a CPU even when a GPU is available on the device, and the like.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content, in conjunction with a declarative language for generating digital content assets and the like may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the common behavior, over a plurality of operating systems. In embodiments, the code execution engine may govern execution of the digital content asset to provide the consistent user experience. A code execution engine may perform a runtime check of the operating system of a device on which the code execution engine is executing the digital content asset control code and adjust a sequence of instructions, a selection of instructions, and the like based on a result of the runtime operating system check. In embodiments, a code execution engine of the system for creating, sharing and managing digital content may ensure consistent user experience of the digital content asset by, for example, decoupling the declarative language from a target system-specific resources, such as a device and/or operating system specific rendering engine and the like. In embodiments, rendering actions, as may be represented by a rendering layer of functionality, may be encoded into the declarative language used in the visual editing environment to generate a digital content asset and the like. In this way, a digital content asset may perform the same behavior on different devices, different operating systems, and combinations thereof. In embodiments, the code execution engine may function similarly to a gaming engine in that the combination of the declarative language with the code execution engine may define behaviors at a rendering level of the digital content asset, such graphic drawing primitives, and the like. In embodiments, this combination may facilitate coding the users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior at the time that the digital content asset control code is generated through use of the declarative language in the visual editing environment. In this way, the code execution engine, which may function similarly to a gaming engine for this specific aspect, may do the underlying work of making the behavior consistent, without a developer having to adjust the use of the declarative language for a digital content asset for each target device and/or operating system that may need to be controlled to generate the desired behavior.

In embodiments, a code execution engine that works cooperatively with a declarative language for digital content asset creating and editing may govern execution of the code for a consistent user experience, such as is described herein, across a plurality of mobile operating systems, including, without limitation operating systems such as IOS™, ANDROID™, WINDOWS™, and the like. Mobile operating systems may include their own look and feel, including how fundamental user interactions are performed. Governing execution of code across mobile operating systems may include use of some device-specific declarative language so that, while the user experience may not be the same across

mobile operating systems, a user of an IOS™ based device may experience the digital content asset with a look and feel that a user experiences when using other mobile applications on the device. In this way, a consistent user experience may be tailored to each individual mobile operating system so that the digital content asset may effectively appear to have the same behavior, while the underlying user interface and mobile operating system native controls may be preserved for each mobile operating system. This may involve, for example, executing portions of the declarative language consistently on each mobile operating system and executing other portions of the declarative language according to the mobile operating system. In embodiments, declarative language statements for rotational behavior of a digital content asset may be consistent across mobile operating systems, whereas declarative language statements for user interface elements for controlling the rotation may be operating system specific.

In embodiments, a code execution engine of a system that uses a declarative language for digital content asset creation and editing may govern execution of the code for a consistent user experience across a plurality of computer operating systems, including, without limitation operating systems such as MAC™, LINUX™, WINDOWS™, and the like. Likewise, the code execution engine may govern execution of the code for a consistent user experience in deployments that include combinations of mobile operating systems (e.g., an IPHONE™) and a computer (e.g., a WINDOWS™ LAPTOP). In embodiments, a combination of the code execution engine and the visual editing environment may facilitate this consistent user experience across mobile, computer, and other operating systems, such as by enabling creation, delivery and editing of the digital content asset during runtime (e.g., when the code execution engine is executing digital content asset control code, and the like). In embodiments, LLVM compilation may facilitate generation of operating system-specific sets of compiled code that may include a portion of the declarative language representation of the digital content asset and a portion of the code execution engine to perform one or more digital content asset control actions consistently across a plurality of operating systems, including mobile operating systems, computer operating systems, and the like.

In embodiments, the code execution engine may enable control of network layer interactions for the digital content asset. In embodiments, the code execution engine may be structured with a layer that may facilitate controlling network layer interactions, such as a network layer of the code execution engine, and the like. In embodiments, the code execution engine may gain control of network layer interactions for the digital content asset via network layer interaction control statements that may be available to a developer in the declarative language during editing. The visual editing environment may make available to a developer and/or other user of the environment network layer interaction declarative language statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, network layer interactions and the like may be controlled.

In embodiments, a system for creating, sharing and managing digital content that may further enable digital content asset generation and control via a declarative language may further comprise a LLVM debugging tool. The LLVM debugging tool may facilitate debugging compiled code for the digital content asset, and the like.

In embodiments, the system for creating, sharing and managing digital content that cooperatively uses a declara-

tive language for generating a digital content asset via a visual editing environment and executed the digital content asset with a code execution engine may further include functions that facilitate orchestrating components, events, response to triggers and the like for a digital content asset. In embodiments, orchestrating may include without limitation automated arrangement, coordination, and management of components, events and the like. Such orchestrating may facilitate control of aspects of the system to accomplish goals of the system, such as just-in-time compilation through use of an LLVM-based compiler architecture and the like. In embodiments, orchestrating functionality may be enabled by a plug-in capability of the system, where an orchestrating capability may be plugged-in to the system. In embodiments, the plug-in capability of the system may be a JAVASCRIPT™ compatible plug-in system.

In embodiments, a system for creating, sharing and managing digital content is depicted in FIG. 31. The system 3100 may comprise a visual editing environment 3102 that may enable a developer 3018 and the like creating, editing, and delivery of a digital content asset 3110 with declarative language 3104. The system 3100 may further comprise a code execution engine 3008 that may facilitate hardware resource utilization and thermal optimization for CPUs, GPU, and the like of different devices, such as end user devices 3012 on which a digital content asset 3110 operates. The system 3100 may be adapted so that the same declarative code 3104 and code execution engine 3008 may be used during visual editing, such as with the visual editing environment 3102 and at runtime 3108. The visual editing environment 3102 and runtime 3108 may both be compiled with an LLVM compiler 3112. The code execution engine 3008 may also control interactions, such as network layer interactions, communication protocol interactions, browser interactions, network middleware and the like. The system 3100 may further comprise an orchestration facility 3014 for orchestrating components, events, and the like associated with the digital content asset 3110. The orchestration facility 3014 may further comprise a plug-in system 3020. The system 3100 may further facilitate consistent user experience across a plurality of different operating systems 3022, including mobile, computer and other operating systems.

In embodiments, a system for creating, sharing and managing digital content may include a visual editing environment, a code execution engine, a declarative language, and a gaming engine capability. The visual editing environment may enable a developer to create and edit code controlling a digital content asset in the declarative language. In embodiments, the visual editing environment may enable a developer to create and edit declarative language code controlling a digital content asset. In embodiments, the code execution engine may operate in the visual editing environment, such as on the created declarative language code to control execution of hardware elements, such as a hardware infrastructure element that enables utilization of the digital content asset. In embodiments, the code execution engine may include and utilize a gaming engine capability to facilitate execution of the declarative language. In embodiments, the code execution engine may utilize the gaming engine's capability to execute the declarative language to further control aspects of the digital content asset including, without limitation a behavior and a state of the digital content asset, and the like.

In embodiments, a system that combines the visual editing environment and a code execution engine with a gaming engine capability to enable at least creating and editing declarative language code for controlling a digital content

assess may include the visual editing environment interacting with the code execution engine during, for example creation of digital content assets. The visual editing environment interacting with the code execution engine may further engage the gaming engine capability of the code execution engine for creation of digital content assets. In embodiments, a user, such as a developer interacting with the visual editing environment may, such as through graphical manipulation of a digital content asset, developing code using statements of the declarative language and the like, may effectively be producing declarative language code that the code execution engine, and optionally the gaming engine capability of the code execution engine responds to by executing it to cause visual effects in the visual editing environment, and the like. The code execution engine, such as through its gaming engine capability may be configured to respond to the produced declarative language code with data and graphic manipulation functions as may be expressed in the declarative language. In embodiments, the declarative language, the code execution engine and the gaming engine capability of the code execution engine that are used during editing of the digital content asset in the visual editing environment may also be used at runtime. In embodiments, the visual editing environment and the code execution engine are constructed from this declarative language. In embodiments, the visual editing environment may make the declarative language available to a developer to generate digital content assets and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of a state, such as a state of a digital content asset that may be expressed in a declarative language. The gaming engine capability may recognize an explicit state of an object, such as a digital content asset and respond based on the context of the object in which the explicit state is expressed. In embodiments, an explicit state expression may be recognized by the gaming engine capability as declaring different properties or values thereof for an object such as a digital content asset, and the like. The gaming engine capability may, for example, recognize that a state in a painting application includes that a pen is selected with a particular color. The code execution engine, such as through the gaming engine capability may render the result of the user's use of the pen and any changes to the pen as a result, such as changing orientation, rotation, and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of an inheritance parameter, such as an inheritance parameter of a digital content asset that may be expressed in a declarative language. The gaming engine capability may recognize expression of an inheritance parameter in the declarative language and respond thereto based on, for example, how the value of the inherited parameter impacts an object, such as an instance of digital content asset being executed by the code execution engine. In embodiments, a declarative language may support creating a sub class that may represent an inheritance parameter operation. The gaming engine capability may determine that when an object is of a sub-class, then it may process parameters of a parent class for the object, effecting enabling an inheritance parameter for the object, such as an instance of a digital content asset and the like. The gaming engine capability of the code execution engine may cause an object, such as an object in a sub-class to operate the same as the parent object. In embodiments, the gaming engine capability may cause an instance of a digital

content asset that is contained in another object to perform actions, such as scaling, rotation, and the like by processing the action properties of the parent object through enabling inheritance of such parameters, and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of an animation feature, such as an animation feature of a digital content asset that may be expressed in a declarative language. The gaming engine capability may recognize expression of an animation feature in the declarative language and respond thereto based on, for example, how the expressed animation feature impacts an object, such as an instance of digital content asset being executed by the code execution engine. A gaming engine capability of the code execution engine may handle an animation feature that is expressed in the declarative language through hardware acceleration of at least a portion of the animation expressed. The gaming engine capability may perform the animations and rendering of transitions of the digital content asset, such as property changes and the like that may be expressed in an animation feature statement and/or due to an impact of the animation on the digital content asset, such as on an instance of a digital content asset being rendered in a user interface, and the like. The gaming engine capability may enable animation features expressed in the declarative language, such as speech animation, procedural animation, skeletal animation, facial animation, 3D animation, and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of a simulation feature, such as a simulation feature of a digital content asset that may be expressed in a declarative language. The gaming engine capability may recognize expression of a simulation feature in the declarative language and respond thereto based on, for example, how the expressed simulation feature impacts an object, such as an instance of digital content asset being executed by the code execution engine. A gaming engine capability of the code execution engine may handle a simulation feature that is expressed in the declarative language through hardware acceleration of at least a portion of the simulation expressed. The gaming engine capability may perform the simulations and rendering of transitions of the digital content asset, such as property changes and the like that may be expressed in a simulation feature statement and/or due to an impact of the simulation of the digital content asset, such as on an instance of a digital content asset being rendered in a user interface, and the like. The gaming engine capability may enable simulation features of a digital content asset, such as an instance of a digital object expressed in the declarative language, such as speech simulation, skeletal simulation, facial simulation, and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of a 3D geometric behavior, such as a 3D geometric behavior of a digital content asset that may be expressed in a declarative language. The gaming engine capability may recognize expression of a 3D geometric behavior in the declarative language and respond thereto based on, for example, how the expressed 3D geometric behavior impacts an object, such as an instance of digital content asset being executed by the code execution engine. A gaming engine capability of the code execution engine may handle a 3D geometric behavior that is expressed in the declarative language through hardware acceleration of at least a portion of the simulation

expressed. In embodiments, the gaming engine capability of the code execution engine may be utilized by the visual editing environment to facilitate rendering three-dimensional visual effects, handling three-dimensional objects, and geometric parameters, such as 3D geometric parameters of objects, such as 3D digital content assets and the like. In 5
embodiments, the gaming engine capability may be embodied in an animation engine portion of the code execution engine. In embodiments, 3D geometric behavior expressed through geometric behavioral statements of the declarative language may facilitate a gaming engine capability of the code execution engine applying rules of physics and geometry on the digital objects for which the geometric behavioral statements of the declarative language are expressed. In 10
embodiments, the gaming engine capability may facilitate 3D geometric behavior for geometrically nested elements (e.g., by use of a scene tree/scene graph of the declarative language and the like as described herein), 3D geometric behavior based on a point of view that may be expressed in the declarative language, and the like.

In embodiments, a gaming engine capability of a code execution engine of a system for creating, sharing and managing digital content may enable handling of a shader functionality, such as shader loading parameters for utilizing a digital content asset in different hardware devices. In 25
embodiments, the shader loading parameters may be expressed in a declarative language. The gaming engine capability may recognize expression of shader loading parameters in the declarative language and respond thereto based on, for example, how the expressed shader loading parameters may impact utilization of an object, such as an instance of digital content asset on different hardware devices. A gaming engine capability of the code execution engine may handle shader loading parameters that are expressed in the declarative language through hardware acceleration, such as through use of GPUs on different hardware devices, and the like. Handling shader loading parameters may be responsive to pixel-handling capacity of a display screen of different hardware devices. In 30
embodiments, recognition of the pixel-handling capacity of a display screen for a hardware device on which the digital content asset is targeted to be utilized may impact how shader loading parameters are handled. The gaming engine capability may adjust how shader loading parameters, including any such parameters that are expressed and/or derived from a shader loading-related expression in the declarative language, are handled, including how they are applied to different devices based on, for example, the pixel-handling capacity of a display screen of the different devices.

In embodiments, a system that uses declarative language to create, and at least edit a digital content asset and uses a code execution engine with a gaming engine capability to enable creation-time and editing-time consistent behavior across different devices may include the visual editing environment interacting with the code execution engine during, for example creation of digital content asset management code, as described herein with declarative language being processed by the gaming engine capability and the like. In 45
embodiments, a user interacting with the visual editing environment may effectively be producing code that portions of the code execution engine, such a gaming engine capability of the code execution engine respond to by causing visual effects in the visual editing environment, the runtime environment, and the like. A gaming engine capability of the code execution engine may be configured to respond to the declarative language code with graphic

manipulation functions for each different device based on, for example, device identification information, such as the type of operating system that the different devices are executing, availability and type of GPU, pixel-handling capacity of the device display, and the like.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content may use a declarative language during editing and the same language during runtime. The use of the same language may facilitate use of a gaming engine capability during editing, runtime, and the like for processing objects, such as digital content asset. This may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the same behavior, over a plurality of devices. 15
In embodiments, the gaming engine capability may govern how the declarative language is executed to provide the consistent user experience. A consistent user experience may include a look and feel of a user interface, a speed of control of a digital content asset, interaction elements in the user interface, and the like. In embodiments, a gaming engine capability may take a declarative language statement for a digital content asset and convert it into a set of pixel manipulation actions that ensure a consistent user experience across different operating systems. This may take the form of graphic primitives, and the like to generate a consistent visual element for each device and operating system. 20

In embodiments, a gaming engine capability enabled code execution engine of the system for creating, sharing and managing digital content may ensure simultaneous user experience of the same behavior of the digital content asset by, for example, decoupling the declarative language from target system-specific resources, such as a device and/or operating system specific rendering engine and the like. In 25
embodiments, rendering actions, as may be represented by a gaming engine capability, may be encoded into the declarative language used in the visual editing environment to generate digital content asset control instances and the like. In this way, a digital content asset may perform the same behavior on different devices, different operating systems, and combinations thereof. In embodiments, the code execution engine may function similarly to a gaming engine through the use of the gaming engine capability in that the combination of the declarative language, and optionally runtime code output therefrom with the gaming engine capability may define behaviors at a pixel-rendering level of the digital content asset, such as 3D movements, and the like. In 30
embodiments, through use of a gaming engine capability in this combination a user in a visual editing environment may code the users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior with the declarative language at the time that the digital content asset is created, edited and the like. In this way, the gaming engine capability of the code execution engine, which may function similarly to a gaming engine for this specific aspect, may do the underlying work of making the behavior consistent, without a developer having to consider how any target device and/or operating system may need to be controlled to generate the desired behavior. 40
50

In embodiments, a code execution engine of a system for creating, sharing and managing digital content may process a declarative language with a gaming engine capability to control utilization of hardware resources of a plurality of different devices, such as CPUs, GPUs and the like. Utilization of, for example, CPUs of some of the different devices, such as hardware endpoint devices and the like may 65

be controlled to facilitate users of different devices experiencing the same behavior. In embodiments, the gaming engine capability of the code execution engine may operate within the visual editing environment during creation, editing, and the like by processing declarative language statements. With the same code execution engine operating during visual editing and during runtime, and the same declarative language being utilized by the visual editing environment and the code execution engine, visual editing may result in generating code, such as digital content asset control code that can control utilization of a CPU and/or GPU, such as by generating code execution control statements. Code execution control statements in a declarative language may include hardware resource utilization statements that a gaming engine capability and the like may process to directly control utilization of different device hardware resources, such as a CPU, GPU, and the like. In embodiments, a language used in the visual editing environment, such as a declarative language, may include hardware resource utilization statements that the code execution engine via the gaming engine capability and the like may execute or that may affect how the code execution engine executes code, such as executing a graphic function with a CPU even when the gaming engine capability determines that a GPU is available on the device, and the like.

In embodiments, the code execution engine of the system for creating, sharing and managing digital content may further control utilization of hardware resources for different aspects of hardware performance, including thermal performance, battery management, and the like. The code execution engine may rely on its gaming engine capability to help control utilization based on aspects such as thermal performance, and the like. A declarative language used to program digital content assets and the like may include statements that facilitate managing execution on target devices to optimize hardware aspects, such as thermal performance and the like. In embodiments, the declarative language may provide access to instruction-level execution power and thermal performance information for different devices. Device-specific instances of a compiled digital content asset, for example, may be represented at the instruction-level so that the impact on at least one of thermal and power performance may be determined for each instruction that may be executed by the code execution engine on the devices. The digital content asset control code created and/or edited, such as by a developer using the declarative language in the visual editing environment can be analyzed based on the power and/or thermal impact of each corresponding device-specific instruction. The result of this analysis may be a measure of the thermal and/or power (e.g., battery demand) impact on the device so that the impact may be controlled. In embodiments, the analysis of the digital content asset control code that the code execution engine may execute may suggest specific code and/or execution control of that code, such as a specific sequence of instructions, a rate of execution of instructions, use of a GPU, use of a CPU, use of a gaming engine capability, and the like that may reduce or optimize thermal performance of the device. In embodiments, optimizing thermal performance for a hardware resource of one or more different devices for which utilization may be controlled, such as a CPU, a GPU and the like may be based on computation of a thermal impact of executing a digital content asset control code set by a code execution engine. This thermal impact computation may include CPU utilization (e.g., execution rate and the like), GPU utilization, memory utilization, and the like and may be determined by the thermal impact of instruc-

tions, such as CPU instructions generated from execution sequences of the gaming engine capability by a compiler, and the like. In embodiments, the thermal impact computation may include compiled instructions generated from a code execution engine performing the digital content asset control code on the device.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content, in conjunction with a gaming engine capability for processing declarative language used for generating digital content assets and the like may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the common behavior, over a plurality of operating systems. In embodiments, the code execution engine may govern execution of the digital content asset to provide the consistent user experience. In embodiments, a code execution engine with a gaming engine capability may ensure consistent user experience of the digital content asset by, for example, decoupling the declarative language from a target system-specific resources, such as a device and/or operating system specific rendering engine and the like. In embodiments, rendering actions may be encoded into the declarative language used in the visual editing environment to generate a digital content asset and processed by the gaming engine capability for each type of operating system. In this way, a digital content asset may perform the same behavior on different devices, different operating systems, and combinations thereof. In embodiments, the code execution engine's gaming engine capability may function similarly to a gaming engine in that the combination of the declarative language with the gaming engine capability may define behaviors at a pixel-rendering level of the digital content asset, such graphic drawing primitives, and the like. In embodiments, this combination may facilitate coding the users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior at the time that the digital content asset control code is generated through use of the declarative language in the visual editing environment. In this way, the code execution engine, through the gaming engine capability may do the underlying work of making the behavior consistent across different operating systems, without a developer having to adjust the use of the declarative language for a digital content asset for each target operating system that may need to be controlled to generate the desired behavior.

In embodiments, a code execution engine that works cooperatively with a declarative language for digital content asset creating and editing may govern gaming engine capability operation for a consistent user experience across a plurality of mobile operating systems, including, without limitation operating systems such as IOS™, ANDROID™, WINDOWS™, and the like.

In embodiments, a code execution engine of a system that uses a declarative language for digital content asset creation and editing may govern execution of the code for a consistent user experience across a plurality of computer operating systems, including, without limitation operating systems such as MAC™, LINUX™, WINDOWS™, and the like. Likewise, the code execution engine may govern execution of the code for a consistent user experience in deployments that include combinations of mobile operating systems (e.g., an IPHONE™) and a computer (e.g., a WINDOWS™ LAPTOP). In embodiments, a combination of the code execution engine and the visual editing environment may facilitate this consistent user experience across mobile, computer, and other operating systems, such as by enabling creation, delivery and editing of the digital content asset

during runtime (e.g., when the code execution engine is executing digital content asset control code, and the like).

In embodiments, a code execution engine equipped with a gaming engine capability may enable control of network layer interactions for the digital content asset. In embodiments, the code execution engine may be structured with a layer that may facilitate controlling network layer interactions, such as a network layer of the code execution engine, and the like. This network layer of the code execution engine may be combined with gaming engine capabilities to facilitate processing declarative language network interaction statements and the like. In embodiments, the code execution engine may gain control of network layer interactions for the digital content asset via network layer interaction control statements that may be available to a developer in the declarative language during editing. The visual editing environment may make available to a developer and/or other user of the environment network layer interaction declarative language statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, and optionally when executed by a gaming engine capability of the code execution engine, network layer interactions and the like may be controlled.

In embodiments, a code execution engine, such as a code execution engine with gaming engine capability as described herein, may gain control of browser interactions for the digital content asset via browser interaction control statements that may be part of digital content asset editing in the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment browser interaction statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, and optionally by the gaming engine capability of the code execution engine, browser interactions and the like may be controlled. In embodiments, browser interactions that may be controlled may include Comet interactions, HTTP streaming interactions, Ajax push interactions, reverse Ajax interactions, Secure Socket interactions, and HTTP server push interactions, and the like. Browser interactions that may be controlled may facilitate browser interactions of one or more different devices that may be connected via a network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the system for creating, sharing and managing digital content that may include a gaming engine capability-enabled code execution engine may further include functions that facilitate orchestrating components, events, response to triggers and the like for a digital content asset. In embodiments, orchestrating functionality may be enabled by a plug-in capability of the system, where an orchestrating capability may be plugged-in to the system. In embodiments, the plug-in capability of the system may be a JAVASCRIPT™ compatible plug-in system.

In embodiments, the visual editing environment, and runtime code produced from editing, creation and other actions of a user of the visual editing environment may be written in the declarative language and may be bound to the code execution engine, such as through compilation, such as with an LLVM-based compiler architecture and the like. In embodiments, the code execution engine may comprise a C++ engine. The code execution engine may perform execution of C++ code; the code execution engine may be coded in C++; the gaming engine capability of the code execution engine may be C++-based; the code execution engine may

facilitate execution of C++ code as well as other code types, such as the declarative language type described herein, and the like.

In embodiments, a system for creating, sharing and managing digital content is depicted in FIG. 32. The system 3200 may comprise a visual editing environment 3102 that may enable a developer 3018 and the like creating, editing, and delivery of a digital content asset 3110 with declarative language 3104. The system 3200 may further comprise a code execution engine 3208 that may facilitate hardware resource utilization and thermal optimization for CPUs, GPU, and the like of different devices, such as end user devices 3012 on which a digital content asset 3110 operates. The code execution engine 3208 may be adapted to provide gaming engine capability 3210 that may handle digital object state, inheritance, animation features, simulation features, 3D geometric behaviors, shader loading parameters, and the like. The system 3200 may be adapted so that the same declarative code 3104 and code execution engine 3208 may be used during visual editing, such as with the visual editing environment 3102 and at runtime 3108. The visual editing environment 3102 and runtime 3108 may both be compiled with an LLVM compiler 3112. The code execution engine 3208 may also control interactions, such as network layer interactions, communication protocol interactions, browser interactions, network middleware and the like. The system 3200 may further comprise an orchestration facility 3014 for orchestrating components, events, and the like associated with the digital content asset 3110. The orchestration facility 3014 may further comprise a plug-in system 3020. The system 3200 may further facilitate consistent user experience across a plurality of different operating systems 3022, including mobile, computer and other operating systems.

In embodiments, a system for creating, sharing and managing digital content may include a visual editing environment, a code execution engine, and a domain-specific declarative language. The system may enable a user, such as a developer to create and edit code controlling a digital content asset using the domain-specific declarative language. The developer may use a visual editing environment of the system to create and edit the digital content asset controlling code with the domain-specific declarative language. The domain-specific declarative language may be used to generate a script for operating the digital content asset in a computing environment, such as a host computing environment. In embodiments, the script may be specified according to which the digital content asset is serialized in the host computing environment. In embodiments, the script may be specified according to which the digital content asset is de-serialized in the host computing environment. In embodiments, the code execution engine may operate on the code controlling the digital content asset generated in the visual editing environment to, for example, control execution of the code to enable utilization of the digital content asset, and the like. The code execution engine may operate cooperatively with the visual editing environment to facilitate controlling the digital content asset while the domain-specific declarative language is being used to create and edit the digital content asset controlling code. In embodiments, the visual editing environment and the code execution engine may enable creation, delivery and editing of the digital content asset. The visual editing environment and the code execution engine may work cooperatively to enable a plurality of versions of runtime code, such as a compiled output of the digital content asset control code created and/or edited in the visual editing environment using the

domain-specific declarative language. In embodiments, the visual editing environment may be written using the domain-specific declarative language as is used to create and/or edit the plurality of runtime versions of the digital content asset.

In embodiments, the system may support producing different types of runtime versions, such as preview versions, portal versions, and the like. A preview runtime version generated with the system may enable a user to preview one or more behaviors of the digital content asset. A preview version may be a limited functionality version that, when the runtime is executed, exhibits a portion of the behaviors of the digital content asset based on, for example, a subset of available statements of the domain-specific declarative language. A preview runtime version may be suitable for use by a preview viewer executing on a computing system, such as a server computing system, and the like. In embodiments, a preview viewer may communicate with the visual editing environment of the system to facilitate previewing behaviors of a digital content asset for which control code is being generated from the domain-specific declarative language. The generated code may be shared with different runtime versions, such as the preview runtime version and the portal runtime version. The preview runtime version may be compiled with an LLVM compiler and the like.

A portal runtime version generated with the system described herein may enable a user to use the digital content asset, such as on a plurality of different devices and the like. A runtime portal version may be shared among users of different devices. A runtime portal version may be accessed by users, such as users of different devices, different operating systems, and the like. A runtime portal version may be configured as a container that enables publication and consumption of the digital content asset across endpoint devices, platforms, and the like. In embodiments, a portal viewer may access the configured runtime portal version and process the content of the container thereby facilitating use of the published digital content asset runtime version. In embodiments, a portal viewer may communicate with the visual editing environment of the system to facilitate viewing behaviors of a digital content asset for which control code is being generated from the domain-specific declarative language. The generated code may be shared with different runtime versions, such as the portal runtime version. The portal runtime version may be compiled with an LLVM compiler and the like.

In embodiments, a system for creating, sharing and managing digital content may facilitate serializing a digital content asset script. Serializing a digital content asset script may enable running the script within a host computing environment without the need for compiling. Serializing may allow access to the digital content asset control code generated with the domain-specific declarative language to operate the digital content asset, such as an active object without the need for compiling. In embodiments, serializing a digital content asset script may allow editing and loading of digital content, such as a digital content asset and the like at runtime into the code execution engine. In embodiments, the domain-specific declarative language may be statically compiled for execution and may be dynamically modified at runtime without compilation. In embodiments, code, such as a script for a digital content asset may be serialized by a serialization engine 112. In embodiments, the domain-specific declarative language may be compiled for execution and may alternatively be executed as a serialized description (e.g., at a text level).

In embodiments, the system for creating, sharing and managing digital content, such as digital content assets, may facilitate serialization of scripts based on a domain-specific declarative language that may be used to describe a visual editing environment through with the domain-specific declarative language may be used to generate digital content asset control code and the like. Serialization may occur by converting tokens, such as words and operators of the domain-specific declarative language, into bytecodes. Serialization may facilitate producing bytecodes that may be associated with literals, such as strings, numbers, object names and the like. An associated literal and the like may be stored with a corresponding bytecode for serialization. In embodiments, a literal may be stored following a bytecode and literal length to facilitate smaller and faster transmission than similar actions involving parsing and the like without serialization. In embodiments, serialized digital content, such as a digital content asset and/or a script for a digital content asset as described herein may include object event logic. In embodiments, such object event logic within a serialized digital content asset and the like may be constrained to consist of a list of parameterized methods. In embodiments, when object event logic for a serialized digital content asset and the like are constrained to the list of parameterized methods, states for the serialized digital content asset may be specified in the domain-specific declarative language. Specified states in the domain-specific declarative language may enable conditional logic for the digital content asset by a combination of the parameterized methods and state-based code execution.

In embodiments, a script for a digital content asset may be de-serialized through use of elements of the system for creating, sharing and managing digital content. A script may be de-serialized for running within a host computing environment without the need for compiling.

In embodiments, a system that enables generation of a plurality of runtime versions of domain-specific declarative language-based digital content assets may enable creation-time and editing-time consistent behavior across different devices. This may include the visual editing environment interacting with the code execution engine during, for example creation of digital content assets from the domain-specific declarative language, and the like. In embodiments, a user interacting with the visual editing environment may effectively be producing code that the code execution engine responds to by executing it to cause visual effects in the visual editing environment, and the like. In embodiments, the visual editing environment may facilitate the users of the different devices to simultaneously experience the same behavior of the digital content asset through a multi-user synchronization system that may operate as part of the visual editing environment to allow, among other things simultaneous experience of the same behavior of a digital content asset while creating and editing in the visual editing environment, and the like.

In embodiments, the visual editing environment and the code execution engine that enables a plurality of runtime versions of digital content asset control code generated from, for example a domain-specific declarative language, may control utilization of hardware resources, such as CPUs, GPUs and the like of different computing devices. Utilization of, for example, CPUs of some of the different devices, such as hardware endpoint devices and the like may be controlled to facilitate the simultaneous experience of the same behavior.

In embodiments, the code execution engine of the system for creating, sharing and managing digital content that users

a domain-specific declarative language to generate a plurality of runtime versions may further control utilization of hardware resources for different aspects of hardware performance, including thermal performance, battery management, and the like. The code execution engine may control utilization based on aspects such as thermal performance, and the like. A domain-specific declarative language used for digital content asset control code and the like may include statements that facilitate managing execution on target devices to optimize hardware aspects, such as thermal performance and the like. In embodiments, the domain-specific declarative language may provide access to instruction-level execution power and thermal performance information for different devices. By targeting a portion of the plurality of runtime versions toward specific devices, the digital content asset control code may be represented at the device instruction-level so that the impact on at least one of thermal and power performance may be determined for each instruction that may be executed by the code execution engine on the target device. The digital content asset control code created and/or edited, such as by a developer using the declarative language in the visual editing environment can be analyzed based on the power and/or thermal impact of each corresponding device-specific instruction. Actions to facilitate optimization of thermal performance, such as specific sequences of code and the like may be made available to the developer in the visual editing environment, and the like.

In embodiments, a code execution engine of the system for creating, sharing and managing digital content, in conjunction with a domain-specific declarative language used for generating digital content assets and the like may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the common behavior, over a plurality of operating systems. In embodiments, the code execution engine may govern execution of the digital content asset to provide the consistent user experience. In embodiments, the domain-specific declarative language may facilitate coding a target users experience (e.g., how the digital content asset will behave) with respect to the digital content asset behavior at the time that the digital content asset control code is generated through use of the domain-specific declarative language in the visual editing environment. In this way, the code execution engine may do the underlying work of making the behavior consistent across different operating systems, without a developer having to adjust the use of the domain-specific declarative language for a digital content asset for each target operating system that may need to be controlled to generate the desired behavior.

In embodiments, a code execution engine that works cooperatively with a plurality of runtime versions generated from domain-specific declarative language may operation on different operating systems for a consistent user experience across a plurality of mobile operating systems, including, without limitation operating systems such as IOS™, ANDROID™, WINDOWS™, and the like.

In embodiments, a code execution engine of a system that uses a domain-specific declarative language for digital content asset creation and editing may govern execution of the code for a consistent user experience across a plurality of computer operating systems, including, without limitation operating systems such as MAC™, LINUX™, WINDOWS™, and the like. Likewise, the code execution engine, optionally in association with runtime versions generated from a domain-specific declarative language may govern execution of the code for a consistent user experi-

ence in deployments that include combinations of mobile operating systems (e.g., an IPHONE™) and a computer (e.g., a WINDOWS™ LAPTOP).

In embodiments, a code execution engine as described herein that works cooperatively with a plurality of runtime versions of a digital content asset generated from domain-specific declarative language may enable control of network layer interactions for the digital content asset. In embodiments, the code execution engine may gain control of network layer interactions for the digital content asset via network layer interaction control statements that may be available to a developer in the domain-specific declarative language during editing. The visual editing environment may make available to a developer and/or other user of the environment network layer interaction domain-specific declarative language statements that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, network layer interactions and the like may be controlled.

In embodiments, a code execution engine, such as a code execution engine that executes one or more versions of a digital content asset generated from a domain-specific declarative language as described herein, may gain control of browser interactions for the digital content asset via browser interaction control statements that may be part of digital content asset editing in the visual editing environment. The visual editing environment may make available to a developer and/or other user of the environment browser interaction statements of the domain-specific declarative language that may be coded into digital content asset behavior and the like so that when executed by the code execution engine browser interactions and the like may be controlled. In embodiments, browser interactions that may be controlled may include Comet interactions, HTTP streaming interactions, Ajax push interactions, reverse Ajax interactions, secure socket interactions, and HTTP server push interactions, and the like. Browser interactions that may be controlled may facilitate browser interactions of one or more different devices that may be connected via a network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the code execution engine may gain control of networking middleware for the digital content asset via networking middleware control statements of the domain-specific declarative language. The visual editing environment may make available to a developer and/or other user of the environment networking middleware statements of the domain-specific declarative language that may be coded into digital content asset behavior and the like so that when executed by the code execution engine, networking middleware and the like may be controlled. In embodiments, networking middleware that may be controlled may facilitate network interaction and the like of Raknet middleware, a gaming engine, a transport layer interaction, a UDP interaction, a TCP interaction, a 3D rendering engine, a gestural engine, a physics engine, a sound engine, an animation engine, and the like. Networking middleware that may be controlled may facilitate network interactions and the like of one or more different devices that may be connected via a network and may, optionally be rendering the digital content asset in a way that ensures a consistent user interface across the different devices.

In embodiments, the system that enables digital content asset control code creation from domain-specific declarative language and the like may further include functionality that facilitate orchestrating components, events, response to trig-

gers and the like for a digital content asset. In embodiments, orchestrating functionality may be enabled by a plug-in capability of the system, where an orchestrating capability may be plugged-in to the system. In embodiments, the plug-in capability of the system may be a JAVASCRIPT™ compatible plug-in system.

In embodiments, a system for creating, sharing and managing digital content is depicted in FIG. 33. The system 3300 may comprise a visual editing environment 3302 that may enable a developer 3018 and the like creating, editing, and delivery of a digital content asset 3110 with domain-specific declarative language 3304. The system 3300 may further enable generation of a plurality of runtime versions 3308. The system 3300 may further comprise a code execution engine 3008 that may facilitate hardware resource utilization and thermal optimization for CPUs, GPU, and the like of different devices, such as end user devices 3012 on which a digital content asset 3110 operates. The system 3300 may be adapted so that the same domain-specific declarative language 3304 and code execution engine 3008 may be used to program the visual editing environment, and the plurality of runtime version 3308. The visual editing environment 3302 and runtime version 3308 may both be compiled with an LLVM compiler. The code execution engine 3008 may also control interactions, such as network layer interactions, communication protocol interactions, browser interactions, network middleware and the like. The system 3300 may further comprise an orchestration facility 3014 for orchestrating components, events, and the like associated with the digital content asset 3110. The orchestration facility 3014 may further comprise a plug-in system 3020. The system 3300 may further facilitate consistent user experience across a plurality of different operating systems 3022, including mobile, computer and other operating systems.

In embodiments, a system for creating, sharing and managing digital content may facilitate efficiency of artist workflow, such as by moving 3D generation into a GPU at runtime. The system may include a visual editing environment, a texture map processing engine, and a 2D-to-3D code generator. In embodiments, the visual editing environment may enable a developer to create and edit code controlling a digital content asset in a declarative language. The developer may specify aspects of the digital content asset such as color, texture and the like for a plurality of layers of an object of a digital content asset. In embodiments, the developer may specify the aspects as described herein in a 2D editing environment that facilitates specifying the aspects for each layer of the digital content asset. In embodiments, the layers of the digital content asset may be 2D layers that may be combined to generate a 3D version of the digital content asset. In embodiments, the texture map processing engine may facilitate processing color or texture information, compressing the color or texture information for the plurality of layers into a texture map data structure. The texture map data structure may represent at least color, texture and the like of the layers of the digital content asset. In embodiments, the 2D-to-3D code generator may apply texture data structure processing operations, such as vertex operations, pixel shading operations, and the like at runtime. By applying the texture map data structure operations at runtime, the code generator may project the object of the digital content asset in 3D at runtime.

In embodiments, the 2D-to-3D code generator may use a generative kernel language when applying texture map data structure processing operations. When a GPU of a hardware environment in which the 2D-to-3D code generator is oper-

ating is available, the code generator may use the GPU at runtime to facilitate projecting the object in 3D.

In embodiments, the system that facilitates 3D projection at runtime with a 2D-to-3D code generation engine may include a gaming engine that may govern behavior of a 3D object, such as 3D object of a digital content asset created with the declarative language. The gaming engine may recognize a 3D object behavior at runtime and respond based on the context of the object in which the 3D behavior is recognized.

In embodiments, the system that facilitates 3D projection at runtime with a 2D-to-3D code generation engine may include a gaming engine that may enable handling of a state, such as a state of a digital content asset expressed in the declarative language. The gaming engine may recognize an explicit state of an object, such as a digital content asset and respond based on the context of the object in which the explicit state is expressed. In embodiments, an explicit state expression may be recognized by the gaming engine as declaring different properties or values thereof for an object such as a digital content asset, and the like.

In embodiments, a gaming engine of the system for creating may enable handling of an inheritance parameter, such as an inheritance parameter of a digital content asset that may be expressed in a declarative language. The gaming engine may recognize expression of an inheritance parameter in the declarative language and respond based thereon, for example, how the value of the inherited parameter impacts an object, such as an instance of digital content asset being executed in a runtime environment. In embodiments, the declarative language may support creating a sub-class that may represent an inheritance parameter operation. The gaming engine may determine that when an object is of a sub-class, it may process parameters of a parent class for the object, enabling an inheritance parameter for the object, such as an instance of a digital content asset and the like. The gaming engine may cause an object, such as an object in a sub-class to operate the same as the parent object due to the gaming engine's ability to handle inheritance parameters expressed in the declarative language and the like.

In embodiments, a gaming engine may enable handling of an animation feature, such as an animation feature of a digital content asset that may be expressed in a declarative language. The gaming engine may recognize expression of an animation feature in the declarative language and respond thereto based on, for example, how the expressed animation feature impacts an object, such as an instance of the digital content asset being executed by the gaming engine. A gaming engine may handle an animation feature that is expressed in the declarative language through hardware acceleration of at least a portion of the animation expressed. The gaming engine may perform the animations and rendering of the digital content asset that may be expressed in an animation feature statement and/or due to an impact of the animation on the digital content asset, such as on an instance of a digital content asset being rendered in a user interface, and the like.

In embodiments, a gaming engine may enable handling of a simulation feature, such as a simulation feature of a digital content asset that may be expressed in a declarative language. The gaming engine may recognize expression of a simulation feature in the declarative language and respond thereto based on, for example, how the expressed simulation feature impacts an object, such as an instance of digital content asset being executed by the code execution engine. A gaming engine may handle a simulation feature that is

expressed in the declarative language through hardware acceleration of at least a portion of the simulation expressed.

In embodiments, a gaming engine may enable handling of a 3D geometric behavior, such as a 3D geometric behavior of a digital content asset that may be expressed in a declarative language. The gaming engine may recognize expression of a 3D geometric behavior in the declarative language and respond game engine features that perform 3D geometric actions. A gaming engine may handle a 3D geometric behavior that is expressed in the declarative language through hardware acceleration of at least a portion of the simulation expressed. In embodiments, 3D geometric behavior expressed through geometric behavioral statements of the declarative language may facilitate a gaming engine applying rules of physics and geometry on the digital objects for which the geometric behavioral statements of the declarative language are expressed.

In embodiments, a gaming engine may enable handling of a shader functionality, such as shader loading parameters for utilizing a digital content asset in different hardware devices. In embodiments, the shader loading parameters may be expressed in a declarative language. The gaming engine may recognize expression of shader loading parameters in the declarative language and respond thereto based on, for example, how the expressed shader loading parameters may impact utilization of an object, such as an instance of digital content asset on different hardware devices. A gaming engine may handle shader loading parameters that are expressed in the declarative language through hardware acceleration, such as through use of GPUs on different hardware devices, and the like. Handling shader loading parameters may be responsive to pixel-handling capacity of a display screen of different hardware devices. In embodiments, recognition of the pixel-handling capacity of a display screen for a hardware device on which the digital content asset is targeted to be utilized may impact how shader loading parameters are handled. The gaming engine may adjust how shader loading parameters, including any such parameters that are expressed and/or derived from a shader loading-related expression in the declarative language, are handled, including how they are applied to different devices based on, for example, the pixel-handling capacity of a display screen of the different devices.

In embodiments, the 2D visual editing environment may enable placement of objects on non-linear planes. The 2D visual editing environment may also enable specifying an effect for a surface of the object, such as a directional light source effect, a shadow effect, a glow effect and the like.

In embodiments, the system may further include a code execution engine that may control utilization of hardware resources, such as CPUs, GPUs and the like of different computing devices. Utilization of, for example, GPUs of some of the different devices, such as hardware endpoint devices and the like may be controlled to facilitate the simultaneous experience of the same behavior.

In embodiments, the code execution engine included in the system may further control utilization of hardware resources for different aspects of hardware performance, including thermal performance, battery management, and the like. The control utilization may include specification and execution of instructions for optimization of thermal performance of a GPU based on execution of the declarative language. A domain-specific declarative language used for digital content asset control code and the like may include statements that facilitate managing execution on target devices to optimize hardware aspects, such as thermal performance and the like.

In embodiments, the code execution engine included with the system may further ensure consistent user experience with the digital content asset, such as a consistent user experience of the common behavior, over a plurality of operating systems. In embodiments, the included code execution engine may govern execution of the digital content asset to provide the consistent user experience. In embodiments, a code execution engine may operate differently on different operating systems for a consistent user experience across a plurality of operating systems, including, without limitation operating systems such as IOS™, ANDROID™, WINDOWS™, MAC™, LINUX™, and the like.

In embodiments, a system for creating, sharing and managing digital content is depicted in FIG. 34. The system 3400 may comprise a visual editing environment 3402 that may enable a developer 3018 and the like to create, edit, and the like a digital content asset. The visual editing environment 3402 may further enable specifying a surface effect of an object placed on a non-linear plane. The visual editing environment 3402 may further facilitate specifying 2D editing of layer-specific color and textures. The visual editing environment 3402 may allow a developer to create a digital content asset 3110 using a declarative language 3404. The system 3400 may further include a texture mapping system 3408 that may facilitate producing a multi-layer texture map data structure from the 2D layer-specific colors and textures specified in the visual editing environment 3402. The system 3400 may further comprise a 2D-to-3D code generator 3420 that may function with a generative kernel language to perform vertex operations, pixel shading operations and object 3D runtime projection. The system 3400 may further comprise a code execution engine 3414 that may facilitate hardware resource utilization and thermal optimization for CPUs, GPU, and the like of different devices, such as end user devices 3012 on which a digital content asset 3110 operates. The system 3400 may further comprise a gaming engine 3412. The code execution engine 3414 may also control interactions, such as network layer interactions, communication protocol interactions, browser interactions, network middleware and the like. The system 3400 may further facilitate consistent user experience across a plurality of different operating systems 3022, including mobile, computer and other operating systems.

In embodiments, a system for creating, sharing and managing digital content may facilitate creating and rendering generative content. The system may perform ingestion of inputs, combine the inputs into an abstract representation of a system that processes the inputs, the combined inputs may be adapted to conform to a simulation environment and simulated to produce output comparable to outputs from the system for processing the inputs. The simulated outputs may be optimized, such as through a simulated annealing process. The system may further synthesize content from the outputs. In embodiments, the system may further clean up the synthesized content, such as with a simulated annealing process. In embodiments, the content may be compressed, transmitted and rendered on different types of computing devices. In embodiments, the system may optionally perform steps including decimation of the content, encoding of the decimated content, storage of the encoded content, querying e.g., of the encoded content, generative kernel language execution by a CPU of the content, apply class-specific decoders to the content, and the like.

In embodiments, executing generative kernel language attempts to recreate the original input. Generative kernel language formatted data is a much smaller representation of

the source input and may comprise parameterized procedural instructions for transmission to rendering devices that may process a series of execution/decompression phases. Fidelity of the representation (e.g., a compression ratio) during decimation may be linked to the level of detail required, with the lower the detail, the simpler the representation.

In embodiments, a generative kernel language may comprise a minimal subset of instructions that facilitate generating outputs in a two-phase geometric digital content process. The generative kernel language may be used in a first phase of execution on the CPU. Secondary phases of the two-phase process may use methods other than generative kernel language on a CPU and or GPU. The secondary phases may be specific to a type of output being generated and therefor may not be restricted to generative kernel language instructions. This second phase may result in a reconstructed output similar to original content being rendered such as data, textures, geometry and the like. In embodiments, generative kernel language may accommodate the process of transforming geometric primitives through a two-phase execution process. The generative kernel language may act as a set of CPU-based parameterized procedural instructions which can be executed to create other CPU and/or GPU instructions or data, which can in turn be used to generate content similar to original content. The created instructions or data may be used in secondary phase to create data, textures, geometry and the like with a CPU or a GPU to a preferred level of rendering fidelity.

In embodiments, a generative content system may perform ingestion and combination of inputs. Inputs may be ingested from many different sources, may be of differing types, differing accuracy, differing precision and the like. Inputs may be ingested independent of a rate of update of the different inputs. Ingesting and combining input may include statistically processing the inputs, temporally filtering the inputs, spatially filtering the inputs, and combining these processed inputs.

In embodiments, a generative content system may use these processed inputs to create abstract instances of classes. These abstract instances of classes may include properties that can connect to one another, such as in a hierarchical graph. In embodiments, the abstract instance of classes may also be spatially and temporally arranged in a data structure. In embodiments, the data structure of the abstract instance of classes may include a data structure partitioning items with up to 4-dimensional axes, such as nested squares forming a quad tree, cubes forming an oct tree, tesseract forming a 'hyper tree', and the like. In embodiments, processing inputs to create abstract instances of classes may be scalable, such as based on the number or volume of inputs, to facilitate, for example batch processing. Results of such batch processing may be combined into a shared representation. In embodiments, the abstract instances of classes may contribute to a graph with any number of nodes. Properties of these nodes may be connected to other nodes (e.g., with comparable properties, and the like) in a tree structure, such as a directed acyclic graph and the like. In embodiments, the nodes may also be partitioned in a plurality of dimensions, such as four dimensions based on the node properties (e.g., time and x, y, z location).

In embodiments, a generative content system may process a graph of classes with a set of class dependent algorithms. The generative content system may iterate over the graph with the class dependent algorithms. The class dependent algorithms may attempt to converge on a set of fitness criteria that may have many arbitrary weighted elements

associated with the graph node hierarchy, classes, node properties, and the like. These algorithms can be specific to a given class or class hierarchy and act at different levels of detail on the graph. In embodiments, these algorithms may alter node properties, such as based on the class and properties of the current node, its parent, children or sibling nodes. These class dependent algorithms may also prune or create new nodes, such as for correcting errors, filling in information, and the like. The degree to which the graph is processed (e.g., the number of iterations and the like) may be adjusted based on a preference for convergence to the fitness criteria versus time/computation. In an example where data sets of road centerlines are provided with a number of lanes, where a minor road centerline joins a major road centerline without a controlled intersection (e.g. no traffic lights), nodes and properties defining a road intersection surface and markings in contact with the ground plane, even stop sign nodes can be added. In embodiments, a densely packed abstract representation may allow faster processing with fewer computing resources, increased centralized computation thereby reducing computation load on end user devices and potentially less network traffic. Additionally, fewer simulation and conformance passes may be needed to be completed before content synthesis operations. Also, synthesis may be performed in a computationally distributed manner (e.g. actually making the vertexes for the 3D representation of the road and signs). In embodiments, a densely packed abstract representation may facilitate efficiently splitting this densely packed representation. In embodiments, the convergence algorithms act on the graph by updating nodes, deleting nodes, creating nodes, updating node properties, deleting node properties, and creating node properties and the like.

In embodiments, once the conformance simulation portion of the generative content system is completed responsive to the fitness criteria, the hierarchical graph may be prepared for the generation of synthetic content from the nodes of the graph. An objective of preparing for generation of synthetic content is to facilitate distributed computation and storage. In embodiments, a partitioning scheme may be applied during this preparation process so that the hierarchical graph may be distributed based on the partitioning scheme. In embodiments, the partitioning scheme may allow multiple sparse data sets to be generated for processing which have edge nodes based on the partition scheme that are marked to be used for other computations. In embodiments, the marked edge nodes do not have results in localized graph operations that use techniques such as procedural rules, simulations, a libraries of templates/lookup tables, AI models and genetic algorithms to create the content at the maximum quality.

Through the use of the partitioning scheme and marking some nodes, it is possible to distribute this process of synthesis because a local copy of the graph may only need to have some nodes marked for content synthesis/generation. In embodiments, these synthetic generative processes, which may be associated with a class, can act on the graph nodes and generate, copy and parametrize a plurality of types of content that may be associated with the nodes.

In embodiments, the generated content may go through a set of specialized conformance processes that may perform clean up and annealing of aspects of the content. Specialized conformance processing may be required at a class level, content type, spatial or temporal level, and the like. In an example, conformance processing may include splitting polygons, adding and aligning vertexes in spatially proximate meshes, such as where a lower resolution ground plane

meets a road surface, normalizing an envelope when combining two generated audio signals, and the like. Specialized conformance processing may be distributed for processing. In embodiments, the specialized conformance processing may be associated with a class can act on one or more content items. These content items can be identified by the type of content, properties of the content itself, the node's properties or the association of other nodes and their properties and content which are connected via hierarchy or partitioning as defined by the process, and the like.

In embodiments, the processed content can go through a set of specialized decimation processes that may be specific to a type of content being processed and a level of detail required to be preserved for rendering, and the like. It is possible to distribute this process as only the current piece of content being processed is required for decimation. These content type-specific decimation processes can act on the processed content for each level of detail required. In embodiments, output of this process may be simplified/quantized/transformed or empty pieces of content for each level of detail based, for example on a set simplification rules that may be encoded in the decimation process, and the like.

In embodiments, the level of detail decimated content may go through content specific encoding processes. These processes may be specific to a type of content and a type of encoding or compression applied. In embodiments, it is possible to distribute these processes because only the specific content item being processed is required for encoding. In embodiments, content-specific coherent batch execution may be preferred due to it is more efficient and may facilitate inter-content techniques like a shared texture atlas or building a statistical model for a lookup table, and the like. In embodiments, these encoding processes output data in a generative kernel language format.

In embodiments, the rendering level of detail generated generative kernel language formatted data may be stored in a multi-dimensional database that may manage volume and spatial occlusion. In embodiments, a generative content system may incorporate the multi-dimensional database. In embodiments, the multi-dimensional database may be distributed over multiple servers.

In embodiments, the multi-dimensional partitioned database may be queried remotely, such as based on an x/y/z position, x, y, z orientation and a time, and the like. The multi-dimensional partitioned database may reply by supplying packets of generative kernel language data ordered from the front of the frustum to the back; thereby allowing rendering to proceed with the most important packets arriving first. These packets may also be supplied with the right level of detail for the proximity to the viewer, taking into account occlusion, and the like. In embodiments, a query of the multi-dimensional database may be more abstract, such as based on the type of rendering required. In embodiments, the content may include generative kernel language holding time series data or simulations of particle interactions, and the like. In embodiments, generative kernel language data packets may be cached by the local client most lazily. This reduces bandwidth and re-computing items which are in view in subsequent frames. This may facilitate achieving high frame rates.

In embodiments, generative kernel language may facilitate local execution in a virtual machine and the like. In embodiments, generative kernel language may comprise a packing and parameterization language that expands into class-specific instructions and data.

In embodiments, a generative content system may support class-specific decoding. A plurality of decoders and processing steps may be applied to the unpacked results from generative kernel language execution. These may be determined by the specific class and can be executed on a CPU or GPU. Examples of class-specific decoders include SE (Spline Extraction) designed to unpack smoothly varying time series data such as joint angles on the CPU, TS (Texture Synthesis) and GML (Generative modelling language) both designed to be executed on the GPU, and the like.

FIG. 35 may depict a flow chart of a generative content approach that may include a method 3500 that may include an ingesting step 3502 that may ingest inputs from a range of sources, having a range of formats, and indicating a range of times. A step of combining 3504 the inputs may follow the ingesting step. A step of building a hierarchical graph of classes 3508 of the inputs may be followed by a step of arranging the classes of the data temporally and spatially in a data structure 3510. The result of step 3510 may be used in a step of applying class-dependent convergence algorithms 3512 to achieve a fitness criteria 3506. The converged data from 3512 may be processed to generate a plurality of portions of the graph with marked synthesis nodes at step 3514. These graph portions may be decimated at step 3518, and encoded at step 3520 to produce generative kernel language formatted data 3522. The generative kernel language formatted data may be executed by a processor at step 3524. Executing the generative kernel language formatted data at step 3524 may produce class-specific rendering instructions and data 3528.

FIG. 36 depicts an example configuration of an application system 3600. In some embodiments, the application system 3600 may be the system 100 shown with respect to FIG. 1A. In the illustrated example, the application system 3600 may be configured to support a multi-user infrastructure and multi-user synchronization. The application system 3600 may include a declaration processor 3602, a scene tree manager 3604, a mirroring processor 3606, a visual editor 3608, an editor broadcaster 3610, and an editor listener 3612. The application system 3600 may include additional components, such as other components discussed with respect to FIG. 1A.

In embodiments, the declaration processor 3602 is configured to linearly process a first declaration. The declaration may recite any set of items selected, including objects, relationships, properties, behaviors, and combinations thereof.

In embodiments, the scene tree manager 3604 may be configured to manage an instantiation of objects, object relationships, properties, and behaviors when embodied in an instantiation of a scene tree. In embodiments, the scene tree manager 3604 updates the instantiation of the scene tree when a user of many user provides user input to change the instantiation from a user computing system 3620 via a communication network 3630. For example, the user may provide input to change an object, an object relationship, a property, and/or a behavior. In response to the user input, the scene tree manager 3604 may update the instantiation to reflect the change to the object, the object relationship, the property, and/or the behavior.

In embodiments, the mirroring processor 3606 enables mirroring the first declaration to a second declaration. In embodiments, the mirroring processor 3606 is configured to cause generation of an instantiation of objects in the instantiation of the scene tree that mirror the set of items in the first declaration from the declaration processor. In embodiments,

the mirroring processor **3606** is further configured to cause mirroring, in whole or in part, of the scene tree in a second declaration.

In embodiments, the visual editor **3608** is configured to receive user input indicating changes to instantiation of the scene tree and to change the instantiation of the scene tree based on the user input via the scene tree manager.

In embodiments, the editor broadcaster **3610** is configured to obtain changes from the visual editor and to cause the mirroring processor to transform a correspondingly changed part of the instantiation of the scene tree into a change declaration. In embodiments, the editor broadcaster **3610** is configured broadcast the change declaration embedded in a change message over the network **3630** to the user computer systems **3620**.

In embodiments, the editor listener **3612** may be in communication with the communication network **3630**. The editor listener **3612** may be configured to receive and process any change messages from any other user computer system **3620** of the user computer systems **3620** by causing the mirroring processor **3612** to transform a declaration embedded in received change message into a corresponding change in the instantiation of the scene tree.

In embodiments, the system **3600** may further implement one or more clocks to synchronize operation across the user computer systems. For example, the system may implement a frame clock that defines consecutive frames of the engine, whereby the systems may synchronize the operations to the frame clock.

As used herein, a “computer process” may refer to the performance of a described function in a computer using computer hardware (such as a processor, field-programmable gate array or other electronic combinatorial logic, or similar device), which may be operating under control of software or firmware or a combination of any of these or operating outside control of any of the foregoing. All or part of the described function may be performed by active or passive electronic components, such as transistors or resistors. In using the term “computer process,” a schedulable entity, or operation of a computer program or a part thereof is not necessarily required; although, in some embodiments, a computer process may be implemented by such a schedulable entity, or operation of a computer program or a part thereof. Furthermore, unless the context otherwise requires, a “process” may be implemented using more than one processor or more than one (single- or multi-processor) computer.

As used herein, an “operating system” is an environment of a host computer system, which may be a conventional operating system, but alternatively may be any computing environment that can cause execution of computer instructions, receive input and provide output, including, for example, a web browser, a Java runtime environment, a Microsoft common language runtime environment, etc.

As used herein, a “simulator” in an application system environment engine is an object-oriented system for managing the scheduling of asynchronous and synchronous behaviors of objects, relationships between objects, and changes to properties of objects.

A first computer system is in communication “in real time” over a network with a second computer system when information is exchanged bidirectionally between the computer systems according to scheduling determined by the first and second computer systems and without dependency on polling.

As used herein, “Linear logic” may refer to a set of procedures that can be executed strictly from left to right, or

top to bottom, without looping or branching, and that can be defined with zero, one, or more parameters.

Detailed embodiments of the present disclosure are disclosed herein; however, it is to be understood that the disclosed embodiments are merely exemplary of the disclosure, which may be embodied in various forms. Therefore, specific structural and functional details disclosed herein are not to be interpreted as limiting, but merely as a basis for the claims and as a representative basis for teaching one skilled in the art to variously employ the present disclosure in virtually any appropriately detailed structure.

The terms “a” or “an,” as used herein, are defined as one or more than one. The term “another,” as used herein, is defined as at least a second or more. The terms “including” and/or “having”, as used herein, are defined as comprising (i.e., open transition).

While only a few embodiments of the present disclosure have been shown and described, it will be obvious to those skilled in the art that many changes and modifications may be made thereunto without departing from the spirit and scope of the present disclosure as described in the following claims. All patent applications and patents, both foreign and domestic, and all other publications referenced herein are incorporated herein in their entireties to the full extent permitted by law.

The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software, program codes, and/or instructions on a processor. The present disclosure may be implemented as a method on the machine, as a system or apparatus as part of or in relation to the machine, or as a computer program product embodied in a computer readable medium executing on one or more of the machines. In embodiments, the processor may be part of a server, cloud server, client, network infrastructure, mobile computing platform, stationary computing platform, or another computing platform. A processor may be any kind of computational or processing device capable of executing program instructions, codes, binary instructions and the like. The processor may be or may include a signal processor, digital processor, embedded processor, microprocessor or any variant such as a co-processor (math co-processor, graphic co-processor, communication co-processor and the like) and the like that may directly or indirectly facilitate execution of program code or program instructions stored thereon. In addition, the processor may enable execution of multiple programs, threads, and codes. The threads may be executed simultaneously to enhance the performance of the processor and to facilitate simultaneous operations of the application. By way of implementation, methods, program codes, program instructions and the like described herein may be implemented in one or more thread. The thread may spawn other threads that may have assigned priorities associated with them; the processor may execute these threads based on priority or any other order based on instructions provided in the program code. The processor, or any machine utilizing one, may include non-transitory memory that stores methods, codes, instructions and programs as described herein and elsewhere. The processor may access a non-transitory storage medium through an interface that may store methods, codes, and instructions as described herein and elsewhere. The storage medium associated with the processor for storing methods, programs, codes, program instructions or other type of instructions capable of being executed by the computing or processing device may include but may not be limited to one or more of a CD-ROM, DVD, memory, hard disk, flash drive, RAM, ROM, cache and the like.

A processor may include one or more cores that may enhance speed and performance of a multiprocessor. In embodiments, the process may be a dual core processor, quad core processors, other chip-level multiprocessor and the like that combine two or more independent cores (called a die).

The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software on a server, client, firewall, gateway, hub, router, or other such computer and/or networking hardware. The software program may be associated with a server that may include a file server, print server, domain server, internet server, intranet server, cloud server, and other variants such as secondary server, host server, distributed server and the like. The server may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other servers, clients, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the server. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the server.

The server may provide an interface to other devices including, without limitation, clients, other servers, printers, database servers, print servers, file servers, communication servers, distributed servers, social networks, and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the server through an interface may include at least one storage medium capable of storing methods, programs, code and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

The software program may be associated with a client that may include a file client, print client, domain client, internet client, intranet client and other variants such as secondary client, host client, distributed client and the like. The client may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other clients, servers, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the client. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the client.

The client may provide an interface to other devices including, without limitation, servers, other clients, printers, database servers, print servers, file servers, communication servers, distributed servers and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the client through an interface may include at least one storage medium capable of storing methods, programs, applications, code and/or instructions. A central repository may provide program instructions to be

executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

The methods and systems described herein may be deployed in part or in whole through network infrastructures. The network infrastructure may include elements such as computing devices, servers, routers, hubs, firewalls, clients, personal computers, communication devices, routing devices and other active and passive devices, modules and/or components as known in the art. The computing and/or non-computing device(s) associated with the network infrastructure may include, apart from other components, a storage medium such as flash memory, buffer, stack, RAM, ROM and the like. The processes, methods, program codes, instructions described herein and elsewhere may be executed by one or more of the network infrastructural elements. The methods and systems described herein may be adapted for use with any kind of private, community, or hybrid cloud computing network or cloud computing environment, including those which involve features of software as a service (SaaS), platform as a service (PaaS), and/or infrastructure as a service (IaaS).

The methods, program codes, and instructions described herein and elsewhere may be implemented on a cellular network having multiple cells. The cellular network may either be frequency division multiple access (FDMA) network or code division multiple access (CDMA) network. The cellular network may include mobile devices, cell sites, base stations, repeaters, antennas, towers, and the like. The cell network may be a GSM, GPRS, 3G, EVDO, mesh, or other networks types.

The methods, program codes, and instructions described herein and elsewhere may be implemented on or through mobile devices. The mobile devices may include navigation devices, cell phones, mobile phones, mobile personal digital assistants, laptops, palmtops, netbooks, pagers, electronic books readers, music players and the like. These devices may include, apart from other components, a storage medium such as a flash memory, buffer, RAM, ROM and one or more computing devices. The computing devices associated with mobile devices may be enabled to execute program codes, methods, and instructions stored thereon. Alternatively, the mobile devices may be configured to execute instructions in collaboration with other devices. The mobile devices may communicate with base stations interfaced with servers and configured to execute program codes. The mobile devices may communicate on a peer-to-peer network, mesh network, or other communications network. The program code may be stored on the storage medium associated with the server and executed by a computing device embedded within the server. The base station may include a computing device and a storage medium. The storage device may store program codes and instructions executed by the computing devices associated with the base station.

The computer software, program codes, and/or instructions may be stored and/or accessed on machine readable media that may include: computer components, devices, and recording media that retain digital data used for computing for some interval of time; semiconductor storage known as random access memory (RAM); mass storage typically for more permanent storage, such as optical discs, forms of magnetic storage like hard disks, tapes, drums, cards and other types; processor registers, cache memory, volatile memory, non-volatile memory; optical storage such as CD, DVD; removable media such as flash memory (e.g., USB sticks or keys), floppy disks, magnetic tape, paper tape,

punch cards, standalone RAM disks, Zip drives, removable mass storage, off-line, and the like; other computer memory such as dynamic memory, static memory, read/write storage, mutable storage, read only, random access, sequential access, location addressable, file addressable, content addressable, network attached storage, storage area network, bar codes, magnetic ink, and the like.

The methods and systems described herein may transform physical and/or intangible items from one state to another. The methods and systems described herein may also transform data representing physical and/or intangible items from one state to another.

The elements described and depicted herein, including in flow charts and block diagrams throughout the figures, imply logical boundaries between the elements. However, according to software or hardware engineering practices, the depicted elements and the functions thereof may be implemented on machines through computer executable media having a processor capable of executing program instructions stored thereon as a monolithic software structure, as standalone software modules, or as modules that employ external routines, code, services, and so forth, or any combination of these, and all such implementations may be within the scope of the present disclosure. Examples of such machines may include, but may not be limited to, personal digital assistants, laptops, personal computers, mobile phones, other handheld computing devices, medical equipment, wired or wireless communication devices, transducers, chips, calculators, satellites, tablet PCs, electronic books, gadgets, electronic devices, devices having artificial intelligence, computing devices, networking equipment, servers, routers and the like. Furthermore, the elements depicted in the flow chart and block diagrams or any other logical component may be implemented on a machine capable of executing program instructions. Thus, while the foregoing drawings and descriptions set forth functional aspects of the disclosed systems, no particular arrangement of software for implementing these functional aspects should be inferred from these descriptions unless explicitly stated or otherwise clear from the context. Similarly, it will be appreciated that the various steps identified and described above may be varied, and that the order of steps may be adapted to particular applications of the techniques disclosed herein. All such variations and modifications are intended to fall within the scope of this disclosure. As such, the depiction and/or description of an order for various steps should not be understood to require a particular order of execution for those steps, unless required by a particular application, or explicitly stated or otherwise clear from the context.

The methods and/or processes described above, and steps associated therewith, may be realized in hardware, software or any combination of hardware and software suitable for a particular application. The hardware may include a general-purpose computer and/or dedicated computing device or specific computing device or particular aspect or component of a specific computing device. The processes may be realized in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable device, along with internal and/or external memory. The processes may also, or instead, be embodied in an application specific integrated circuit, a programmable gate array, programmable array logic, or any other device or combination of devices that may be configured to process electronic signals. It will further be appreciated that one or more of the processes may be realized as a computer executable code capable of being executed on a machine-readable medium.

The computer executable code may be created using a structured programming language such as C, an object-oriented programming language such as C++, or any other high-level or low-level programming language (including assembly languages, hardware description languages, and database programming languages and technologies) that may be stored, compiled or interpreted to run on one of the above devices, as well as heterogeneous combinations of processors, processor architectures, or combinations of different hardware and software, or any other machine capable of executing program instructions.

Thus, in one aspect, methods described above and combinations thereof may be embodied in computer executable code that, when executing on one or more computing devices, performs the steps thereof. In another aspect, the methods may be embodied in systems that perform the steps thereof, and may be distributed across devices in a number of ways, or all of the functionality may be integrated into a dedicated, standalone device or other hardware. In another aspect, the means for performing the steps associated with the processes described above may include any of the hardware and/or software described above. All such permutations and combinations are intended to fall within the scope of the present disclosure.

While the disclosure has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the present disclosure is not to be limited by the foregoing examples, but is to be understood in the broadest sense allowable by law.

The use of the terms “a” and “an” and “the” and similar referents in the context of describing the disclosure (especially in the context of the following claims) is to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. Recitations of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate the disclosure and does not pose a limitation on the scope of the disclosure unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the disclosure. The use of the term “set” includes a group of one or more members.

While the foregoing written description enables one of ordinary skill to make and use what is considered presently to be the best mode thereof, those of ordinary skill will understand and appreciate the existence of variations, combinations, and equivalents of the specific embodiments, method, and examples herein. The disclosure should therefore not be limited by the above described embodiments, method, and examples, but by all embodiments and methods within the scope and spirit of the disclosure.

Any element in a claim that does not explicitly state “means for” performing a specified function, or “step for” performing a specified function, is not to be interpreted as a

“means” or “step” clause as specified in 35 U.S.C. § 112(f). In particular, any use of “step of” in the claims is not intended to invoke the provision of 35 U.S.C. § 112(f).

Persons of ordinary skill in the art may appreciate that numerous design configurations may be possible to enjoy the functional benefits of the inventive systems. Thus, given the wide variety of configurations and arrangements of embodiments of the present invention the scope of the invention is reflected by the breadth of the claims below rather than narrowed by the embodiments described above.

What is claimed is:

1. A system for creating applications, comprising:
one or more processors; and
at least one memory, wherein the one or more processors and the at least one memory implement:
a visual editing environment configured to receive code that defines an application using a declarative language from one or more developers, wherein the code includes a scene tree description;
a code execution engine configured to operate on the code created in the visual editing environment to create a scene tree that defines properties and behaviours of the application based on the scene tree description, wherein the scene tree is a set of objects that collectively define the properties and actions of the application; and
a publisher configured to publish the application, whereby instances of the application may be downloaded to respective client devices, wherein each instance of the application includes a respective instantiation of the scene tree,
wherein the scene tree description is editable via the visual editing environment, such that edits to the scene tree description are updated to the respective instantiations of the scene tree at runtime,
wherein the visual editing environment and the code execution engine enable creation, delivery, and editing of the digital asset during runtime, such that a plurality of end users using different devices can simultaneously experience a same behavior of the digital content asset during its creation and editing, and
wherein the code execution engine governs execution of the code across a plurality of computer operating systems.
2. The system of claim 1, wherein the set of object of the scene tree is a hierarchical map of objects.
3. The system of claim 1, wherein the scene tree specifies respective page layouts of the application.
4. The system of claim 1, wherein the set of objects of the scene tree include one or more variations that manage different states of the application.
5. The system of claim 4, wherein the one or more variations define different languages that are presented in the application.

6. The system of claim 4, wherein the one or more variations define different operating systems that can host the application.

7. The system of claim 4, wherein the one or more variations define different layouts of the application.

8. The system of claim 4, wherein the one or more variations define whether to override objects, properties, or actions of the application.

9. The system of claim 4, wherein the one or more variations define one or more custom states that are specific to the application.

10. The system of claim 4, wherein the one or more variations are defined using variation rules that define different properties of the application based on the different states of the application, wherein each variation rule is either true or false in a given scenario.

11. The system of claim 1, wherein the scene tree description is simultaneously editable by multiple users.

12. The system of claim 1, wherein the code execution engine controls utilization of at least one of a CPU and a GPU of a client device upon which the digital content asset runs.

13. The system of claim 12, wherein control of utilization includes specification and execution of instructions for optimization of thermal performance of the at least one of the CPU and the GPU.

14. The system of claim 1, wherein the code execution engine governs execution of the code across a plurality of operating systems to provide a consistent user experience with the digital content asset.

15. The system of claim 14, wherein the code execution engine governs execution of the code across a plurality of mobile operating systems.

16. The system of claim 15, wherein the mobile operating systems are selected from among IOS®, Android® and Windows® operating systems.

17. The system of claim 1, wherein the code execution engine enables control of network layer interactions for the digital content asset.

18. The system of claim 1, wherein the code execution engine enables control of browser interactions for the digital content asset.

19. The system of claim 18, wherein the browser interactions include Comet interactions, HTTP streaming interactions, Ajax push interactions, reverse Ajax interactions, secure socket interactions, and HTTP server push interactions.

20. The system of claim 1, further comprising a plug-in system for orchestrating components and events for the digital content asset.

21. The system of claim 20, wherein the plug-in system is a Javascript plug-in system.

22. The system of claim 1, wherein the code execution engine is a low-level binary that targets different cross platform device architectures.

* * * * *